Wax

Aydano B. Carleial
Martin E. Hellman
Stanford University
19 July 1974

Packet Radio Temporary Note # 113
SRI Temporary Note #2

BISTABLE BEHAVIOR OF ALOHA-TYPE SYSTEMS

Aydano B. Carleial*
Martin E. Hellman**

* Instituto de Pesquisas Espaciais (INPE), the space research agency of Brazil, and currently at the Department of Electrical Engineering, Stanford University, Stanford, Calif., 94305.
This author's work was partially supported by the U.S. Air Force Office of Scientific Research, under Contract F44620-73-C-0065.

** Department of Electrical Engineering, Stanford University, Stanford, Calif., 94305; and consultant to Stanford Research Institute, Menlo Park, Calif. The work of this author was supported by the Advanced Research Projects Agency of the Department of Defense, under Contract DAHC15-73-C-0187.

ABSTRACT

Packet switching has found widespread application in computer communications because of its ability to efficiently handle high ratios of peak to average data rate. Packet radio is the application of packet switching techniques to radio channels. The resultant multiple access problem requires novel approaches. Such approaches have been developed by others and have primarily been analyzed in steady-state behavior. This paper demonstrates and analyzes an important aspect of the dynamic characteristics of packet radio, namely that of bistable behavior. That is, the system possesses two statistically stable equilibrium points, one in a desirable, low delay region, and the other in an undesirable, high delay region. Since the stability is only statistical in nature the system oscillates between these two points. Even if the resultant steady-state behavior is very poor, this dynamic analysis frequently shows that system performance will be acceptable. This is due to quiet periods (such as at night) which allow the system to recover.

I. Introduction

In this paper we are concerned with communication systems using digital data packets in a multiple-access radio channel. Multiple-access communication problems arise when several sources of information must reach a single destination over a common channel. Examples of this rather common situation are many earth stations simultaneously transmitting to a satellite, or newsmen in a press conference attempting to address the President of the United States. Techniques have been developed to achieve reliable communications in a multiple-access situation. Besides physically separating signals (directive antennas, cables, etc.), one can separate them by partitioning time, or the frequency spectrum, among the transmitters. The resultant time-division multiple access (TDMA) and frequency-division multiple access (FDMA) systems are efficient and practical in most situations. More sophisticated schemes have been devised to obtain higher rates*.

---

*It has recently been shown [2] that the highest communication rates over a multiple-access channel without feedback are achieved through a scheme of superimposed signals, with all transmitters operating all the time and using the entire frequency band. The key to this method is the receiver's ability to reliably decode one of the signals in the presence of interference and noise. After subtracting out this signal, it can proceed with similar decoding steps on the remaining waveform.

Not all multiple-access situations warrant the use of such techniques, however: important considerations often include expandability and flexibility, suitability to time-varying load, simplicity of implementation, time delays, etc.

Computer communications using digital data packets are characterized by short high-rate bursts occurring at random points in time. The statistical nature of the users' demand and the high peak-to-average ratio of data rates make user-assigned TDMA or FDMA inadequate. Consider a system with 400 terminals, of which 100, on the average, are active. Let each of these generate an average of one data packet each 100 seconds, to be sent to the central receiver. Each packet has 1000 bits, and the system's assigned frequency band permits a flow of 10^4 bit/sec into the receiver. Conventional TDMA or FDMA can be used, by assigning time or frequency slots to the terminals, but long delays will be incurred. It takes 40 seconds just to transmit a packet (at $10^4/400 = 25$ bit/sec). At the cost of sign-on procedures and added complexity, one could improve on this performance, using some form of demand assignment, but even then the delay will be at least ten seconds.

Aloha systems [1] were proposed to cope with this type of multiple-access problem and make much shorter delays possible. They are random time-division multiple access (RTDMA) systems, because terminals make use of the channel at random times. There are no user-assigned time slots.

System users generate information at a random rate. The data is formatted into standard-size packets, each carrying the terminal's address, parity-check bits for error detection, and other required information. In an Aloha system, the channel is available to any terminal whenever it has a packet ready for transmission. The entire frequency

Digitized by the Internet Archive
in 2025 with funding from
Amateur Radio Digital Communications, Grant 151

https://archive.org/details/temporary-note-113

band is made available, allowing transmission in a relatively short burst. Provided no other terminal attempts to use the channel during the same time, the communication is successfully accomplished. Sometimes, however, overlapping transmission will occur due to the independent, random origination of packets at the terminals. It is assumed that interference will cause overlapping packets to be lost by the receiver, which therefore does not acknowledge them. The receiver detects (and ignores) damaged packets using the parity-check bits. Each terminal stores its last transmitted packet. If an acknowledgement is not received, the terminal enters the retransmission mode and retransmits the packet, until it gets through to the receiver. Retransmission times are randomized, to prevent recurrent interference among the same terminals. New messages can be generated only when the terminal is cleared. We say that it has then returned to the origination mode.

The channel is not fully utilized, since it is idle part of the time and at times it is subjected to destructive interference. It has been shown [1], nevertheless, that such a system can operate well in practice, provided it is not too heavily loaded. It matches the randomness of user demand in a simple and natural way.

The delay involved in successful transmission of a packet depends on how many times it needs to be retransmitted and how long the terminal must wait to retransmit. A good system should not require too many retransmissions. The average wait time for a retransmission should be short, typically much shorter than the average time between new packet originations. This is consistent with a situation where, at any given time, most of the terminals are in the origination mode. Backlogged packets will normally get through after a short delay, provided there

are only a few terminals in the retransmission mode. If, however, a large number of interfering transmissions occur, the number of retransmission-mode terminals rises and a runaway effect may take place. Due to the high repetition rate of retransmissions, interference becomes so frequent that fewer and fewer transmissions are successful. Still more users are placed in the retransmission mode, more interference takes place, and so on.

Based on the above heuristics, we conjectured that systems would exhibit a bistable behavior. One possible stable condition (with short delays and reasonably good throughput) would have most terminals in the origination mode. The other stable condition would have most terminals in the retransmission mode, obstructing the channel with prevailing interference. Transitions from one condition to the other should occasionally take place, due to statistical fluctuations. Analysis and simulations have confirmed this conjecture. The objective of this paper is to examine bistable behavior and conditions for successful operation of Aloha systems.

We shall demonstrate bistable RTDMA systems of both the synchronous and nonsynchronous type. In a synchronous (or slotted) system, transmissions may start only at precise clock times. A common time reference must be supplied to all terminals and different propagation delays have to be compensated for. The clock period is chosen such that each time slot can accommodate one packet transmission. The analysis of synchronous systems can proceed in discrete time. In a nonsynchronous system, terminals transmit packets without a clock reference. Partially overlapping transmissions interfere and communication capacity is consequently

smaller than in the synchronous case. Simplicity is the great advantage. Synchronous operation may not be feasible in systems with variable propagation delays (e.g., mobile terminals).

We shall model the generation of new packets by an active terminal as a Poisson process for a nonsynchronous system, because this is a good characterization of a thought process and also simplifies analysis. Furthermore, for a large number of users the analysis becomes quite insensitive to the exact distribution of individual terminal traffic generation. The overall traffic generation process tends to a Poisson process. Let the time rate of new packet origination per terminal be denoted by λ_o. Then the average time a user waits before transmitting while in the origination mode is $1/\lambda_o = T_o$.

As previously mentioned, packet generation is inhibited when the terminal is in the retransmission mode. Retransmission of the backlogged packet will then be attempted, also at random points in time*, until successful reception takes place. This artificial process of repetitions has rate λ_r. Practical system designs would tend to have λ_r much greater than λ_o in order to obtain short delays. Bistable behavior results for a range of values under this condition. This model is reasonable if the transmission time for a packet is much smaller than $T_o = 1/\lambda_o$ and $T_r = 1/\lambda_r$. These conditions are satisfied in practice

---

*Although a real system will probably use a bounded randomized wait time, we model this time as an exponential random variable, with mean value $1/\lambda_r$. Again the analysis is not very sensitive to the actual distribution, provided there are many terminals.

for all systems with more than a few terminals.

In the case of synchronous systems it is convenient to use a discrete-time model. An origination-mode terminal will transmit a new packet in any given time slot with probability p_o, whereas retransmissions take place with probability p_r. If τ is the duration of a time slot, p_o and p_r correspond to $\tau\lambda_o$ and $\tau\lambda_r$ in the continuous time model, respectively. The important feature of either characterization is the lack of memory, given the terminal's current mode.

Kleinrock and Lam [3, 4] have independently demonstrated the possibility of bistable behavior in synchronous systems. Their techniques are somewhat different and it is felt that the two approaches complement and add to one another. The reader is referred to [3] and [4] for details.

Exact analysis of slotted systems is simpler and we develop analytic results for this case first. Then we analyze the behavior of unslotted systems with the aid of simulations. A comparison shows that the two types of system behave qualitatively the same. This equivalence is further explored in Section IV, which treats the limiting behavior as the number of users tends to infinity. It is found that, except for a scaling of the data rates, the analysis is almost identical.

II. Synchronous RTDMA Systems

We now proceed with analysis of slotted Aloha systems, as characterized in the last section. Consider N terminals, each emitting packets at random time slots. Transmission is successful if and only if a single terminal attempts communication during the time slot. An unsuccessful attempt puts the terminal temporarily in the retransmission mode.

The channel may be viewed as a discrete-time system with $N+1$ possible states, corresponding to the number $X(t)$ of terminals in the retransmission mode at the beginning of time slot $t = 1, 2, 3, \ldots$. The sequence of system states forms a discrete-time Markov chain with state-transition matrix $P = [p_{nm}]$ where $p_{nm} = P_r[X(t+1) = m \mid X(t) = n]$. In one step, the system may move down one state (a successful retransmission); stay at the same state (no transmissions, or one successful new transmission, or an interference event involving retransmissions only); or move higher in the chain by one or more states (interference event involving one or more fresh packets). We thus find for $n = 0,1,2,\ldots,N$,

$$p_{n,m} = 0 \qquad \text{for } m \leq n-2 \tag{1}$$

$$p_{n,n-1} = (1-p_0)^{N-n} np_r(1-p_r)^{n-1}$$

$$p_{n,n} = (1-p_0)^{N-n}(1-np_r(1-p_r)^{n-1}) + (N-n)p_0(1-p_0)^{N-n-1}(1-p_r)^n$$

$$p_{n,n+1} = (N-n)p_0(1-p_0)^{N-n-1}(1-p_r)^n$$

$$p_{n,m} = \binom{N-n}{m-n} p_0^{m-n}(1-p_0)^{N-m} \qquad \text{for } m \geq n+2$$

For $0 < p_o < 1$ and $0 < p_r < 1$, the Markov chain is irreducible, aperiodic and positive-recurrent, with a vector of long-term state occupation probabilities $\underline{\mu} = (\mu_0, \mu_1, \ldots, \mu_N)^T$ satisfying the equation $\underline{\mu}^T = \underline{\mu}^T P$ [5, pp. 247-253].

The system's expected packet flow per time slot (effective throughput) in a given state n is found to be $f_n = (1-p_o)^{N-n-1}(1-p_r)^{n-1}[(N-n)p_o(1-p_r) + n(1-p_o)p_r]$ for $n = 0,1,2,\ldots,N$. In steady state, the system's throughput is $\bar{f} = \underline{\mu}^T \underline{f} = \sum_{n=0}^{N} \mu_n f_n$, with $\underline{f} = (f_0, f_1, \ldots, f_N)^T$. If $p_o = p_r = p$, the expected packet flow is the same in all states, $f_n = N\,p(1-p)^{N-1} = \bar{f}$. We note here that this throughput is maximized by choosing $p = 1/N$, in which case it converges (from above) to $\bar{f} = e^{-1} \doteq 0.368$ as $N \to \infty$.

Statistics on the transitions from each state are given by the state-transition matrix P. As an indicator of system dynamics, it is interesting to compute the expected drift from state $n = 0,1,2,\ldots,N$. This is simply

$$\bar{a}_n = \sum_{m=0}^{N} (m-n) p_{n,m} \quad . \tag{2}$$

A graph of expected drift versus present state gives a good idea of how the system tends to move in its random walk over the state space. Clearly, $\bar{a}_0$ is positive and $\bar{a}_N$ is negative, so that there is at least one neighborhood of states where the expected drift is near zero and has negative slope. This constitutes an equilibrium point for the system. Since $\bar{a}_n$ changes from positive to negative in such a region, we call it a stable equilibrium point: small excursions result in statistical drifts tending to restore the equilibrium. The system, therefore, lingers

in the region. If $p_o = p_r$, for example, $\bar{a}_n$ is monotone decreasing in n (and approximately linear); there is a single equilibrium point, a stable point. It is indeed verified that states in this neighborhood have the highest occurrence probabilities. See figure 1.

With $p_r > p_o$, as in a practical implementation, the system can exhibit different types of behavior. Figure 2 shows graphs of expected drift and state probabilities for different values of the parameters. Consider the case of figure 2b. Starting from state zero, the system will normally drift toward the first stable point, above which $\bar{a}_n$ becomes negative. At this point the quick dispatch of backlogged packets prevents a rapid rise in state number. Note, however, that the concept of stability is employed only in a statistical sense in the present context. Beyond this stable point lies a second equilibrium region, this one unstable (because $\bar{a}_n$ is changing from negative to positive here). If the system passes this point, it is likely to drift to a still higher, stable region of states, characterized by recurrent interferences, few successful transmissions and long delays. It will usually not return to the more desirable, lower operating point for a long time.

Statistical noise will cause the system to oscillate between these two stable points. The average time spent in one stable region before moving to the other depends on the depth of the "potential well" associated with each region and on the magnitude of the statistical noise in each region (the variance of the distance moved, for example).

Note that the steady-state behavior of a system will depend crucially on the proportion of time spent in each region. For example, if the average time to move from the "good" region to the "bad" region is

one year and the average time to move in the opposite direction is 100 years then, in steady-state only about 1% of the system's time is spent in the good region. However, if the system is lightly loaded at night and is fully utilized for only 8 hours a day, then the mean time between failures due to overloading will be about 1,000 days. When such a failure occurs users will leave the system (probably cursing), thereby allowing it to restart in a good initial state the next day.

We are interested in operating the system in a region of states with reasonably high packet flow and short delay. Over a period of time $[0,T]$, define the average delay D_T (in time slots per packet) to be

$$D_T = \frac{\sum_{t=0}^{T} X(t)}{\sum_{t=0}^{T} \phi(t)} \tag{3}$$

where $X(t)$ is the state and $\phi(t)$ is the number of packets successfully transferred (zero or one) at time slot t. D_T is a random variable whose distribution depends on the path taken by the system's random walk. Letting $\bar{D} = \lim_{T\to\infty} D_T$, we see that

$$\bar{D} = \lim_{T\to\infty} \frac{\left[\sum_{t=0}^{T} X(t)\right]/T}{\left[\sum_{t=0}^{T} \phi(t)\right]/T} \quad \text{and therefore,} \tag{4}$$

$$\bar{D} = \frac{\sum_{n=0}^{N} \mu_n n}{\sum_{n=0}^{N} \mu_n f_n} = \frac{\bar{n}}{\bar{f}} \quad \text{with probability one.} \tag{5}$$

$\bar{D}$ is thus equal to the expected value of the delay for transmission

of a packet, in steady-state*.

$\bar{D}$ can be considerably larger than typical values of D_T obtained while the system remains in low-numbered states (i.e., paths of the random walk that do not visit the higher stable region, in the bistable case). It is important to investigate the possibility of exploring transient operation of the system in the favorable states, in particular by computing statistics on transition times between the two stable regions.

Define the random variable t_{nm} as the time (number of time slots) it takes the system to reach state m for the first time, starting from state n. A method described in [5, pp. 241-246] can be used to determine the characteristic function $\phi_{nm}(u)$ of this random variable and thereby find its distribution. A modified Markov chain is employed where state m is changed into an absorbing state. The other states then become non-recurrent. The analysis gives correct results because, up to the first visit to state m, the modified chain behaves the same as the original chain.

The characteristic function is defined as $\phi_{nm}(u) = E\left(e^{iut_{nm}}\right)$. We take $t_{mm} = 0$ and thus $\phi_{mm}(u) = 1$. For $n \neq m$, we can write

*Radio wave propagation delay was not considered. Only the wait time resulting from the need to retransmit was included. Accordingly, a packet received on the first attempt has delay zero.

$$\phi_{nm}(u) = \sum_{\ell=0}^{N} p_{n\ell}\, E\left(e^{iut_{nm}} \mid \text{first transition is } n \to \ell\right) =$$

$$= p_{nm}\, e^{iul} + \sum_{\ell \neq m} E\left(e^{iu(1+t_{\ell m})}\right) p_{n\ell}$$

$$= e^{iu}\left[p_{nm} + \sum_{\ell \neq m} p_{n\ell}\, \phi_{\ell m}(u)\right] . \qquad (6)$$

If we now define the N-dimensional vectors

$$\underline{\phi}_{(m)} = (\phi_{0m}, \phi_{1m}, \ldots, \phi_{m-1,m}, \phi_{m+1,m}, \ldots, \phi_{Nm})^T \qquad (7)$$

$$\underline{P}_{(m)} = (p_{0m}, \phi_{1m}, \ldots, p_{m-1,m}, p_{m+1,m}, \ldots, p_{Nm})^T$$

and the $N \times N$ matrix $Q_{(m)}$, obtained from P by eliminating its m^{th} row and its m^{th} column, the following equation holds

$$e^{-iu}\, I\underline{\phi}_{(m)} = \underline{P}_{(m)} + Q_{(m)}\underline{\phi}_{(m)} . \qquad (8)$$

The solution

$$\underline{\phi}_{(m)} = \left(e^{-iu}\, I - Q_{(m)}\right)^{-1} \underline{P}_{(m)} \qquad (9)$$

provides the characteristic function of t_{nm} for each initial state $n \neq m$. Note that $P_{(m)}$ and $Q_{(m)}$ depend on the choice of final state. The probability mass function of t_{nm} is found by expressing ϕ_{nm} in the power series form

$$\phi_{nm} = \sum_{t=1}^{\infty} \pi_{nm}(t)\, e^{iut} \qquad (10)$$

We have $Pr\{t_{nm} = t\} = \pi_{nm}(t)$.

Computing the distribution of t_{nm} for a large system may be quite laborious. While all the moments can be found by differentiation

of its characteristic function, the following formula for the mean is obtained directly, by the same method:

$$\begin{bmatrix} \bar{t}_{0m} \\ \cdots\cdots \\ \bar{t}_{m-1,m} \\ \\ \bar{t}_{m+1,m} \\ \cdots\cdots \\ \bar{t}_{Nm} \end{bmatrix} = (I - Q_{(m)})^{-1} \begin{bmatrix} 1 \\ 1 \\ 1 \\ \cdots \\ \cdots \\ 1 \\ 1 \end{bmatrix} \qquad (11)$$

If n is in a stable region of states and m is in the other stable region, $\bar{t}_{nm}$ and $\bar{t}_{mm}$ are good measures of the mean time for transitions from one region to the other.

A computationally straightforward technique for obtaining partial information on the statistics of t_{nm} is also possible. Take the state transition matrix P and modify it by changing m into an absorbing state. Then compute powers of $P_{(m)}$ and observe how the $(n,m)^{th}$ element grows to unity. With a sequence of k matrix-squaring operations, one can compute $P^2_{(m)}$, $P^4_{(m)}$, $P^8_{(m)}$, ..., $P^{2^k}_{(m)}$ and thereby find the probability that state m is reached by the system after 2^k steps or less. Figure 3 shows the result of such a computation for the bistable system of figure 2b, with $n = 0$ and $m = 44$. Assuming T_o to be one minute, we have $\tau = T_o p_o = 0.0075$ min $= 0.45$ sec. The graph shows that the probability of reaching state 44 (in the high delay region) within one hour of operation is about 55%. This is unacceptable system performance. However, any practical implementation of an Aloha system will probably involve many more than 50 users since the advantage over TDMA and FDMA grows with N. It is easier to obtain acceptable system performance for a larger value of N because the

probability of failure within a given time period (e.g., one hour) goes to zero rapidly in N, provided the system is scaled appropriately to handle the larger number of users. This effect is discussed in Section IV.

III. Nonsynchronous RTDMA Systems

A nonsynchronous RTDMA (unslotted Aloha) system cannot be described solely by a Markov chain. Rather, the appropriate stochastic process model is that of a semi-Markov process [6]. If $\{X_i\}_{i=1}^{\infty}$ is a Markov chain, then $X(t)$ is a derived semi-Markov process if

$$X(t) = X_i, \quad \sum_{j=1}^{i-1} T_j \leq t < \sum_{j=1}^{i} T_j \tag{12}$$

where $\{T_j\}_{j=1}^{\infty}$ is a sequence of random variables with the following properties.

1. Given $\{X_i\}_{i=1}^{\infty}$ the random variables $\{T_j\}_{j=1}^{\infty}$ are independent.
2. Given $\{X_i\}_{i=1}^{\infty}$, T_j and T_k have the same distribution if $X_j = X_k$ and $X_{j+1} = X_{k+1}$.

In other words, the length of time the semi-Markov process spends in a particular state depends on that state and the next state of the process, and given those two states the time of residence is independent of all other times of residence and states occupied.

For an unslotted Aloha system let $X(t)$ equal the number of users in the retransmission mode just after the last time the channel was quiet. Figure 4 shows a typical sample function from a process with five users. The pulses indicate transmissions, and the number just over a pulse is the identity (1 - 5) of the user transmitting. It is assumed that any overlap of pulses (interference) causes all affected packets to fail. The height of a pulse is immaterial, and the only reason for the differing heights is to distinguish pulses. The intervals between transitions will be called epochs.

A little thought shows that under these convontions $X(t)$ is a

semi-Markov process. Considerable information can be obtained by considering only the associated discrete-time Markov process

$$X_i = X\left(\sum_{j=1}^{i-1} T_j\right) \quad . \tag{13}$$

Its statistics are completely specified by the initial state X_0 (usually taken as 0) and its state transition matrix $P = [p_{ij}] \quad 0 \leq i,j \leq N$, where N is the number of users. The elements p_{ij} are defined by

$$p_{ij} = \Pr(X_{n+1} = j | X_n = i) \quad . \tag{14}$$

For example, $p_{i,i-1}$ is the probability that during an epoch a user in the retransmission mode transmits and is not interfered with; $p_{i,i+2}$ is the probability that two users in the origination mode interfere with each other and/or with any number of users in the retransmission mode; etc. Except for certain special values of i and j the p_{ij} are difficult to compute analytically. However, they are easily obtained by simulation and can then be used to predict system behavior. This is easier than estimating system behavior directly. Basically, what we have done is to parameterize the estimation problem.

Once the p_{ij} are known within a reasonable accuracy, gross behavior of the system can be inferred from

$$\bar{a}_n = \sum_{j=0}^{N} (j-n)\, p_{nj} \quad . \tag{15}$$

The quantity $\bar{a}_n$ is the expected distance moved in a transition from state n. As for slotted systems, the only stable regions are those in which $\bar{a}_n$ changes sign from positive to negative as n increases.

Figure 5 shows $\bar{a}_n$ versus n for $N = 25$ and six values of $T_o = T_r \equiv T$. Message duration τ was taken to be one second. Two hundred transitions were simulated for each point plotted. This number is large enough for inferring gross behavior but should not be used as an indication of exact behavior. We see that when $T_o = T_r$, $\bar{a}_n$ is a decreasing function of n, just as in the slotted case. Also, as one would expect, for n fixed $\bar{a}_n$ is a decreasing function of T. This is because T is the mean time until the first user transmits, and for N large it is also approximately the mean time between all transmissions which occur in state n. A larger value of T/N therefore indicates a smaller probability of interference, causing the drift to be more toward lower numbered states.

Now let us turn to the more interesting and realistic case where $T_o \gg T_r$. The feedback effect is then present and various behaviors are possible. If the system is conservatively designed then $\bar{a}_n < 0$ for most states, and even when most users are in the retransmission mode the net drift is toward lower numbered states. For example, in Figure 6, when $T_o = 10\ T_r$ we find $\bar{a}_n < 0$ for $i > 1$. Such a system spends most if its time near states 1 and 2. At the opposite extreme, a poorly designed system has $\bar{a}_n > 0$ for most states and spends most of its time in high numbered states. In Figure 6, $T_o = 20\ T_r$ is an example of such a system. It too has a single equilibrium point, in this case between states 24 and 25.

A bistable situation can also exist, as shown in Figure 6 with $T_o = 15\ T_r$. The plot of $\bar{a}_n$ versus n has three crossings of the axis; the first and last are stable equilibrium points while the middle crossing is an unstable equilibrium point.

IV. Poisson Approximation for Large Values of N

When the number of users, N, is large, a simplifying approximation can be made that yields the expected drift values $\{\overline{a}_n\}_{n=0}^{N}$ for both synchronous and nonsynchronous Aloha systems without recourse to simulation.

Considering nonsynchronous systems first, let us establish the duration of a packet transmission, τ, as our unit of time (i.e., $\tau = 1$). If all terminals were in the origination mode the generation of transmissions would be a Poisson process with intensity $\Lambda_o = N\lambda_o$. If they were all in the retransmission mode, the intensity would be $\Lambda_r = N\lambda_r$. If a fraction r of the terminals are in the retransmission mode, then the intensity of the Poisson process is $\Lambda(r) = (1-r)\Lambda_o + r\Lambda_r$. The rate at which users move from the origination mode into the retransmission mode is $(1-r)\Lambda_o[1-\exp(-2\Lambda(r))]$, since $(1-r)\Lambda_o$ is the rate of generation of new transmissions and $[1-\exp(-2\Lambda(r))]$ is the probability that a transmission is not successful. Similarly the rate at which terminals leave the retransmission mode is $r\Lambda_r \exp(-2\Lambda(r))$. The net rate of increase of users in the retransmission mode is the difference between these two quantities and is also equal to $N\,\overline{a(r)}$, where $a(r)$ denotes the time rate of change of r. We now regard r as the system state. It follows that

$$\overline{a(r)} = N^{-1}[(1-r)\Lambda_o - \Lambda(r)\exp(-2\Lambda(r))]. \qquad (16)$$

Note that (16) gives the true expected time rate of change of r (with τ as the unit of time), whereas in Section III we calculated $\overline{a(n)}$, the expected drift per epoch.

If $\Lambda_o = \Lambda_r = \Lambda$, then (16) reduces to

$$\overline{a(r)} = N^{-\frac{1}{2}}[(1-r)\Lambda - \Lambda\exp(-2\Lambda)] \quad , \qquad (17)$$

which is linear in r and crosses the r axis at the stable point $r = 1 - \exp(-2\Lambda)$. If Λ_r is increased, then for reasonable values of Λ_o ($\Lambda_o < 0.5$, to be exact) the stable equilibrium point moves to a lower value of r. See Figure 7. This is desirable since it results in smaller delays. If, however, Λ_r is increased beyond a certain point, two stable equilibria develop. The smaller of these, which we shall denote r_1, continues to move to smaller, more desirable values of r; but the other, denoted r_3, is in the region of high delay. Fortunately, if N is very large then even a small stable region holds the system for a very long time, because the law of large numbers has a chance to set in during the system's random walk on r.

With reference to Figure 7, we see that if $\Lambda_o = 0.2$ and $\Lambda_r = 3.0$ the smaller stable value of r is 0.087. At this point the system throughput is equal to the rate of generation of new packets and is thus equal to $(1-0.087)\Lambda_o = 0.183$. The maximum possible throughput for an unslotted Aloha system is $1/(2e) \doteq 0.184$, and occurs for $\Lambda_o = \Lambda_r = 0.5$. Thus we pay only a slight penalty in throughput, but, as we shall see, gain significantly in delay characteristics. The expected delay for a packet when the system is in a state r is*

$$\overline{D(r)} = \left[\frac{1}{\exp(-2\Lambda(r))} - 1\right]T_r = [\exp(2\Lambda(r)) - 1]\,T_r \qquad (18)$$

*More precisely, $\overline{D(r)}$ is the expected delay given that the system's random walk does not drift away from state r during the packet's transmissions.

since the probability of success on each transmission is $\exp(-2\Lambda(r))$ and each retransmission delay has an average value of $T_r \equiv 1/\Lambda_r$. If $\Lambda_o = 0.2$, $\Lambda_r = 3.0$ and the system is in the lower stable region, the expected delay is $1.43\ T_r = 0.095\ T_o$. For $\Lambda_o = \Lambda_r = 0.5$, the delay is $1.72\ T_o$. We lose only about 0.5% in throughput, but gain a reduction in delay by a factor if 18, a very advantageous tradeoff.

If $\Lambda_o = 0.1845$ and $\Lambda_r = 30$ the lower stable equilibrium point moves to $r = 0.0088$. Under these conditions the throughput is 0.183, the same as above, but the delay is $1.45\ T_r = 0.0089\ T_o$, about one tenth of its previous value. This process can be continued, letting Λ_o approach 0.183 and Λ_r approach infinity. The throughput at the stable region can be kept constant, while the delay tends to zero. In fact, it can be seen that if the throughput r_1 is kept equal to a constant η, then increasing Λ_r by a given factor (say, ten) reduces the delay by exactly the same factor at the lower stable point r_1. This is because at an equilibrium point the throughput η is equal to the rate of generation of new packets $(1-r)\Lambda_o$. Thus $\Lambda(r_1) = (1-r_1)\Lambda_o + r_1\Lambda_r = \eta + r_1\Lambda_r$ and

$$\overline{a(r_1)} = N^{-1}[\eta - (\eta + r_1\Lambda_r)\exp(-2(\eta + r_1\Lambda_r))] = 0 \quad . \qquad (19)$$

In general, if $\overline{a(r)} = 0$ for given values of η and Λ_r then $\overline{a(r/k)} = 0$ for the same value of η provided we now use $\Lambda_r' = k\Lambda_r$. Note that $\Lambda'(r/k)$ is the same as $\Lambda(r)$ even though the stable point moves. From (18) we thus see that the expected delay, measured in terms of T_r, is constant. But T_r is reduced by a factor of k when Λ_r is increased by a factor of k. We thus have

Theorem 1: When operating at a fixed throughput η at a stable equilibrium point, an unslotted Aloha system with an essentially infinite number of users has an expected delay which is inversely proportional to the rate of retransmissions λ_r.

It will be noted that in the above examples where η was held at 0.183 and Λ_r was increased from 3 to 30, the delay did not decrease by exactly a factor of 10. This was due to roundoff error.

The price that is paid for increasing Λ_r is that the unstable equilibrium point, denoted r_2, also moves closer to the origin. In the above examples, with $f_{r_1} = \eta = 0.183$ as Λ_r increases from 3 to 30, r_2 moves from 0.243 to 0.0131. For a finite number N of terminals, there is a limit to how small r_2 can be before the stability of the equilibrium at r_1 becomes threatened. For example, if $r_1 = 0.01$ and $r_2 = 0.02$ and $N = 10^6$ then 10,000 users are in the retransmission mode at the stable equilibrium point r_1, and another 10,000 would have to be placed in the retransmission mode before the unstable point r_2 is reached. This appears to be an acceptable situation. If, however, $N = 1{,}000$ and r_1 and r_2 are still 0.01 and 0.02, only 10 users are in the retransmission mode at the stable point r_1, and only an additional 10 must reach this mode before the system goes unstable. The point r_1 thus appears to be a rather ephemeral stable point.

Turning now to slotted Aloha systems with an essentially infinite number of users, we find, keeping the same notation,

$$\overline{n(r)} = N^{-1}[(1-r)\Lambda_o - \Lambda(r)\exp(-\Lambda(r))] \qquad (20)$$

where $\overline{a(r)}$ is the expected change in r during one time slot (or per unit time since $\tau = 1$). Now $\Lambda_0 = Np_0$ is the average number of transmission attempts in a time slot when $r = 0$ and $\Lambda_r = Np_r$ is the average number when $r = 1$. In obtaining (20) we have used the Poisson approximation to the binomial distribution so that

$$\Pr\left\{\begin{matrix} k_1 \text{ new transmission attempts} \\ \text{and} \\ k_2 \text{ retransmission attempts} \end{matrix}\right\} =$$

$$= \frac{e^{-(1-r)\Lambda_0}\left[(1-r)\Lambda_0\right]^{k_1}}{k_1!} \cdot \frac{e^{-r\Lambda_r}(r\Lambda_r)^{k_2}}{k_2!} . \qquad (21)$$

Comparison of (16) and (20) shows that if a slotted Aloha system has parameters Λ_0 and Λ_r, and if an unslotted system has parameters $\Lambda_0/2$ and $\Lambda_r/2$, then their expected drift curves $\overline{a(r)}$ are identical (within the Poisson approximation) except for an additional factor of two which appears in $\overline{a(r)}$ for the slotted system. This has no real effect on system behavior, and has no effect on the locations of the points r_1, r_2 and r_3. With slight additional reasoning we can prove the following:

Theorem 2: If an unslotted system with an essentially infinite number of users is changed to a slotted system, and the number of users N is doubled, then the equilibrium points r_1, r_2 and r_3 are unmoved and the delay is unaffected.

We thus see that Theorem 1 also applies to slotted systems, and we have a complete equivalence between the two types of system. The only real difference is that unslotted systems have half the capacity of slotted systems.

V. Discussion

Previous analyses of Aloha systems have assumed the traffic generation to be a Poisson process with constant intensity. This tacitly assumes that the average time a terminal waits before retransmitting, T_r, is equal to the average thought time to produce a packet, T_o. Since T_o is typically on the order of one minute, this would result in unacceptably long delays. Any practical system will have a much shorter value of T_r, typically less than one second. This results in Λ_r/Λ_o ratios of 100 and even larger. The results presented in this paper show that such a system is capable of bistable behavior, with one stable point being in a desirable, low-delay region, and the other stable point being in the undesirable, high-delay region. We have given qualitative guidelines concerning system design to ensure that the undesirable stable point is not reached. In particular, most systems will have peak periods during which they are heavily loaded followed by slack periods of low use. For example, the peak period may be eight hours and the slack period may be sixteen hours. If the probability of the system reaching the undesirable region within eight hours is small, then many days will elapse before the system "fails". And even when it fails the following slack period will allow it to return to the desirable region. Thus the prediction that the system spends a long time in the undesirable region is really mathematical fiction.

It may be possible to use a dynamically adaptive strategy to reduce, or completely eliminate even the occasional failures. If T_r were to be increased when the system was approaching the unstable

region, then the position of the unstable region would be moved further from the current operating point, and a margin of safety added. As the system recovered, T_r could again be reduced. The delay would be longer during the recovery period, but at least the delay is allowing greater throughput. If a user is to suffer a given delay it is best for him to interfere with as few other packets as possible.

There are at least two techniques which could be used to control T_r. If there is a central control function which can monitor the channel then it can control the value of T_r through occasional commands to the terminals. Alternatively, a terminal could increase T_r whenever it had to retransmit a packet and decrease T_r whenever a packet was successfully transmitted. The increment sizes would have to be carefully chosen, but it appears that such a strategy could be effective.

We cannot consider exercising a similar control over T_0, since it is our assuption that this parameter, unlike T_r, is a characteristic of the users' thought processes. It would be possible, nonetheless, to send an inhibiting command from the central control to the terminals which would temporarily block all new packet generation (say, by locking keyboards). This would be done whenever some dangerously high state was reached, thereby assuring no further rise in state. Retransmissions would continue, however, and normal operation would resume when some specified lower state was reached. A design tradeoff results between the duration of the outages and their frequency. Note that in all centrally-controlled strategies only an estimate of the current state is available (from the intensity of traffic), but for large N and

$\Lambda_r \gg \Lambda_o$ such an estimate can be quite accurate.

In summary, we have shown that bistable operation is not only possible, but is probably to be found in any practical system. It is hoped that the techniques developed in this paper will provide a basis for designing packet radio systems with smaller delays and less frequent failures.

Acknowledgment

The authors wish to thank Dr. Stanley Fralick of SRI for many valuable discussions. They also thank Mr. Parimi Murthy, currently at the University of Hawaii, for performing the simulations on unslotted Aloha systems.

REFERENCES

1. N. Abramson, "The Aloha System - Another Alternative for Computer Communications", 1970 Fall Joint Computer Conference, pp. 281-285.

2. A. Wyner, "Recent Results in the Shannon Theory", IEEE Trans. on Info. Theory, vol. IT-20, pp. 2-10, Jan. 1974.

3. L. Kleinrock and S. Lam, "On Stability of Packet Switching in a Random Multi-Access Broadcast Channel", Seventh Hawaii Int'l. Conference on System Sciences, Honolulu, Jan. 8-10, 1974, Proc. of the Special Subconference on Computer Nets, pp. 73-77.

4. S. Lam, Packet Switching in a Multi-Access Broadcast Channel with Application to Satellite Communication in a Computer Network, Ph.D. thesis, Computer Science Dept., UCLA, March 1974.

5. E. Parzen, Stochastic Processes, Holden-Day, San Francisco, 1962.

6. R. Pyke, "Markov Renewal Processes: Definitions and Preliminary Properties", Annals of Mathematical Statistics, vol. 32, pp. 1231-1242, 1961.

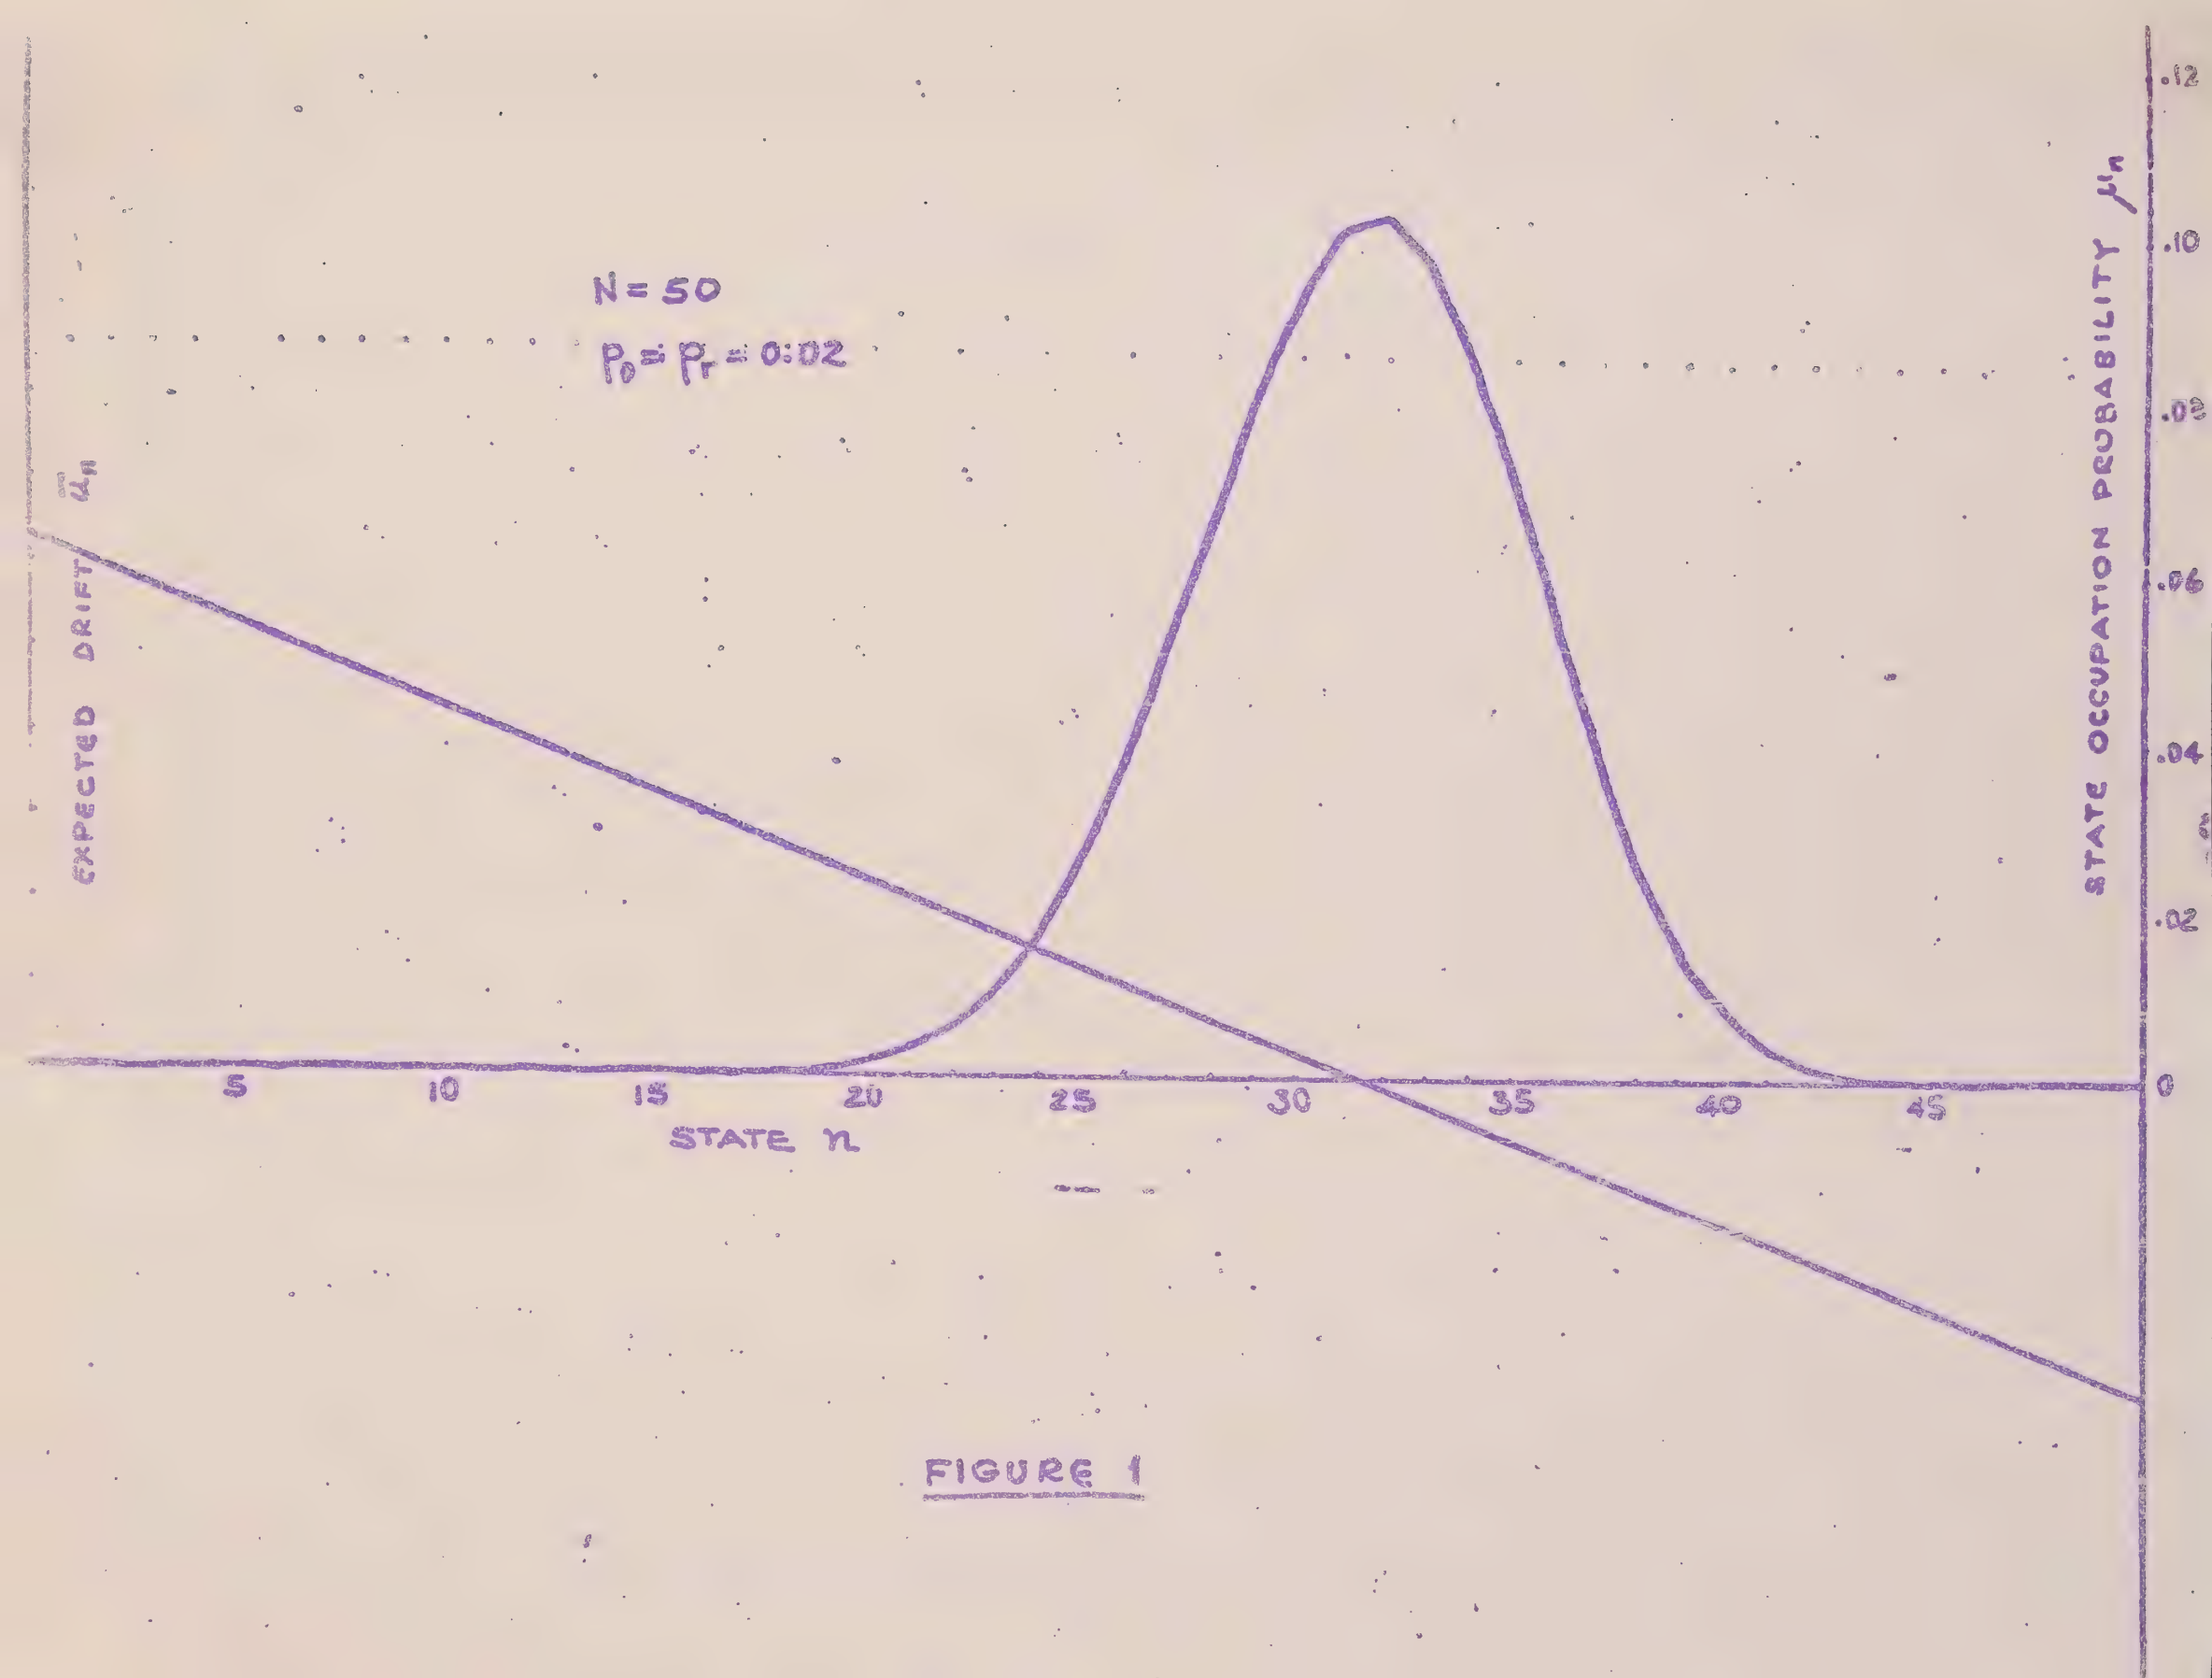
N = 50
$P_0 = P_r = 0.02$
EXPECTED DRIFT μ_n
STATE OCCUPATION PROBABILITY μ'_n
.12
.10
.08
.06
.04
.02
0
5
10
15
20
25
30
35
40
45
STATE n

FIGURE 1

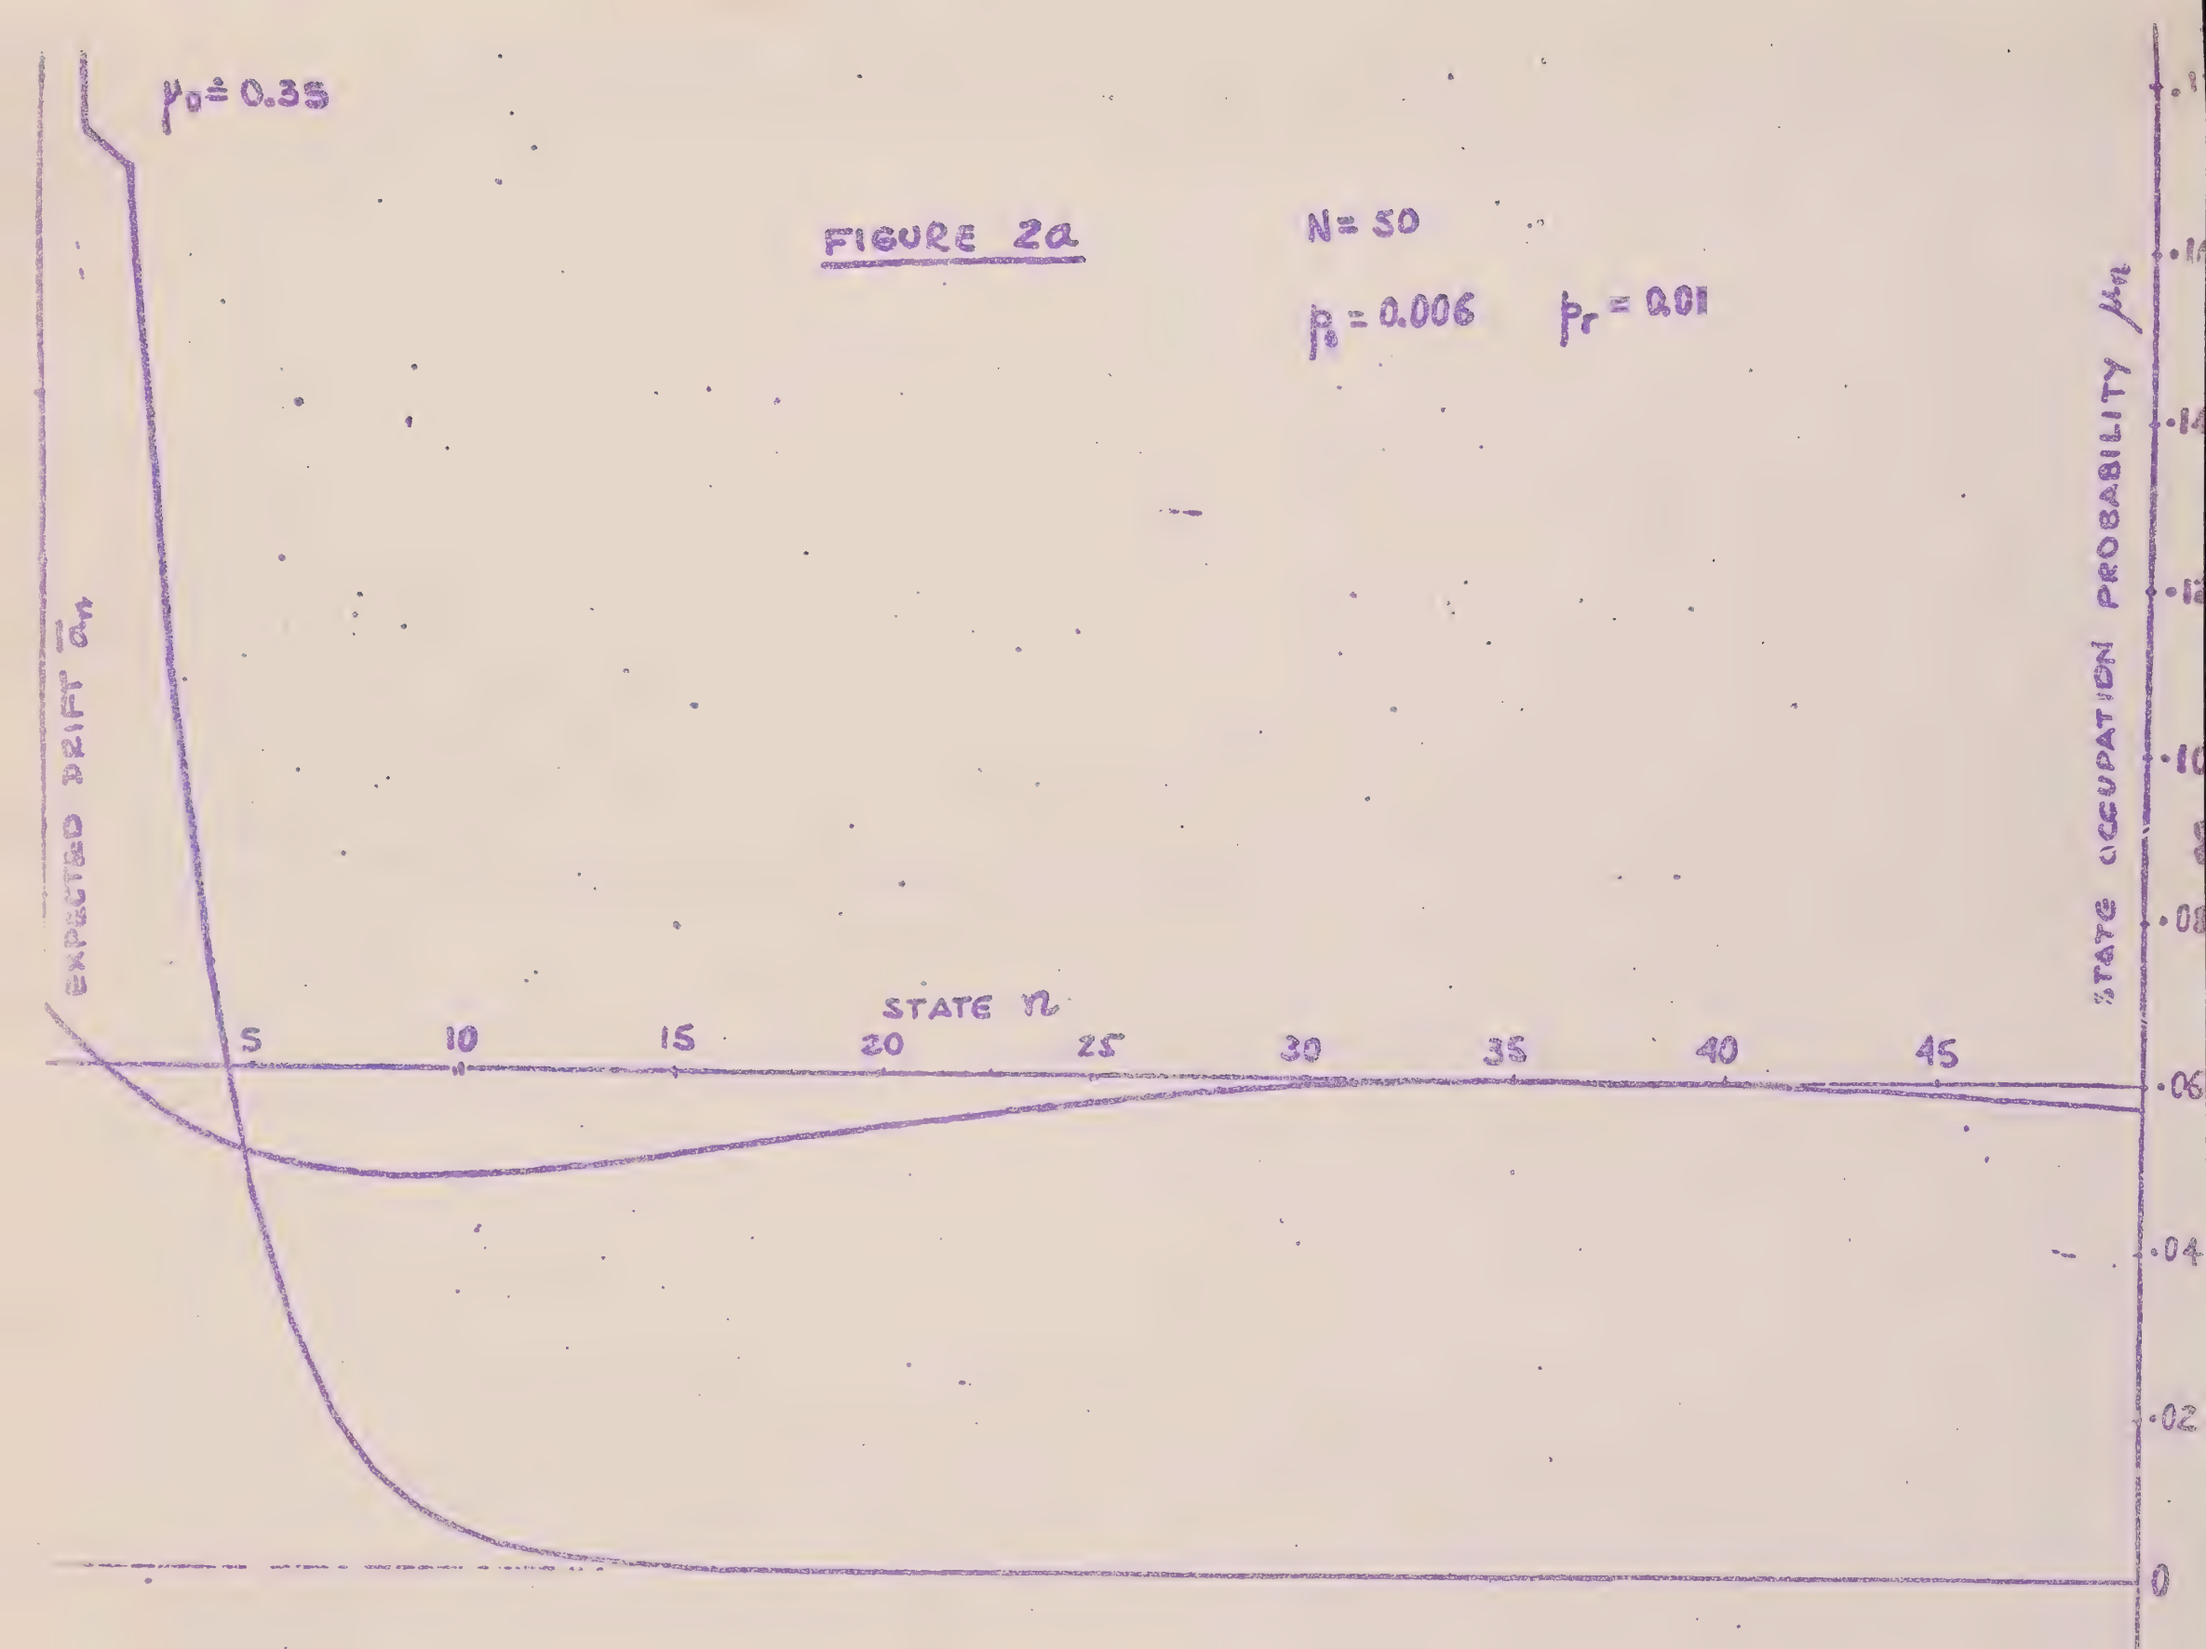
$\mu_0 = 0.35$
FIGURE 2a
N = 50
$p_b = 0.006$
$p_r = 0.01$
EXPECTED DRIFT $\bar{a}_n$
STATE n
5
10
15
20
25
30
35
40
45
STATE OCCUPATION PROBABILITY μ_n
.1
.08
.06
.04
.02
0

FIGURE 2b

N = 50

$P_0 = 0.0075$ $P_r = 0.1$

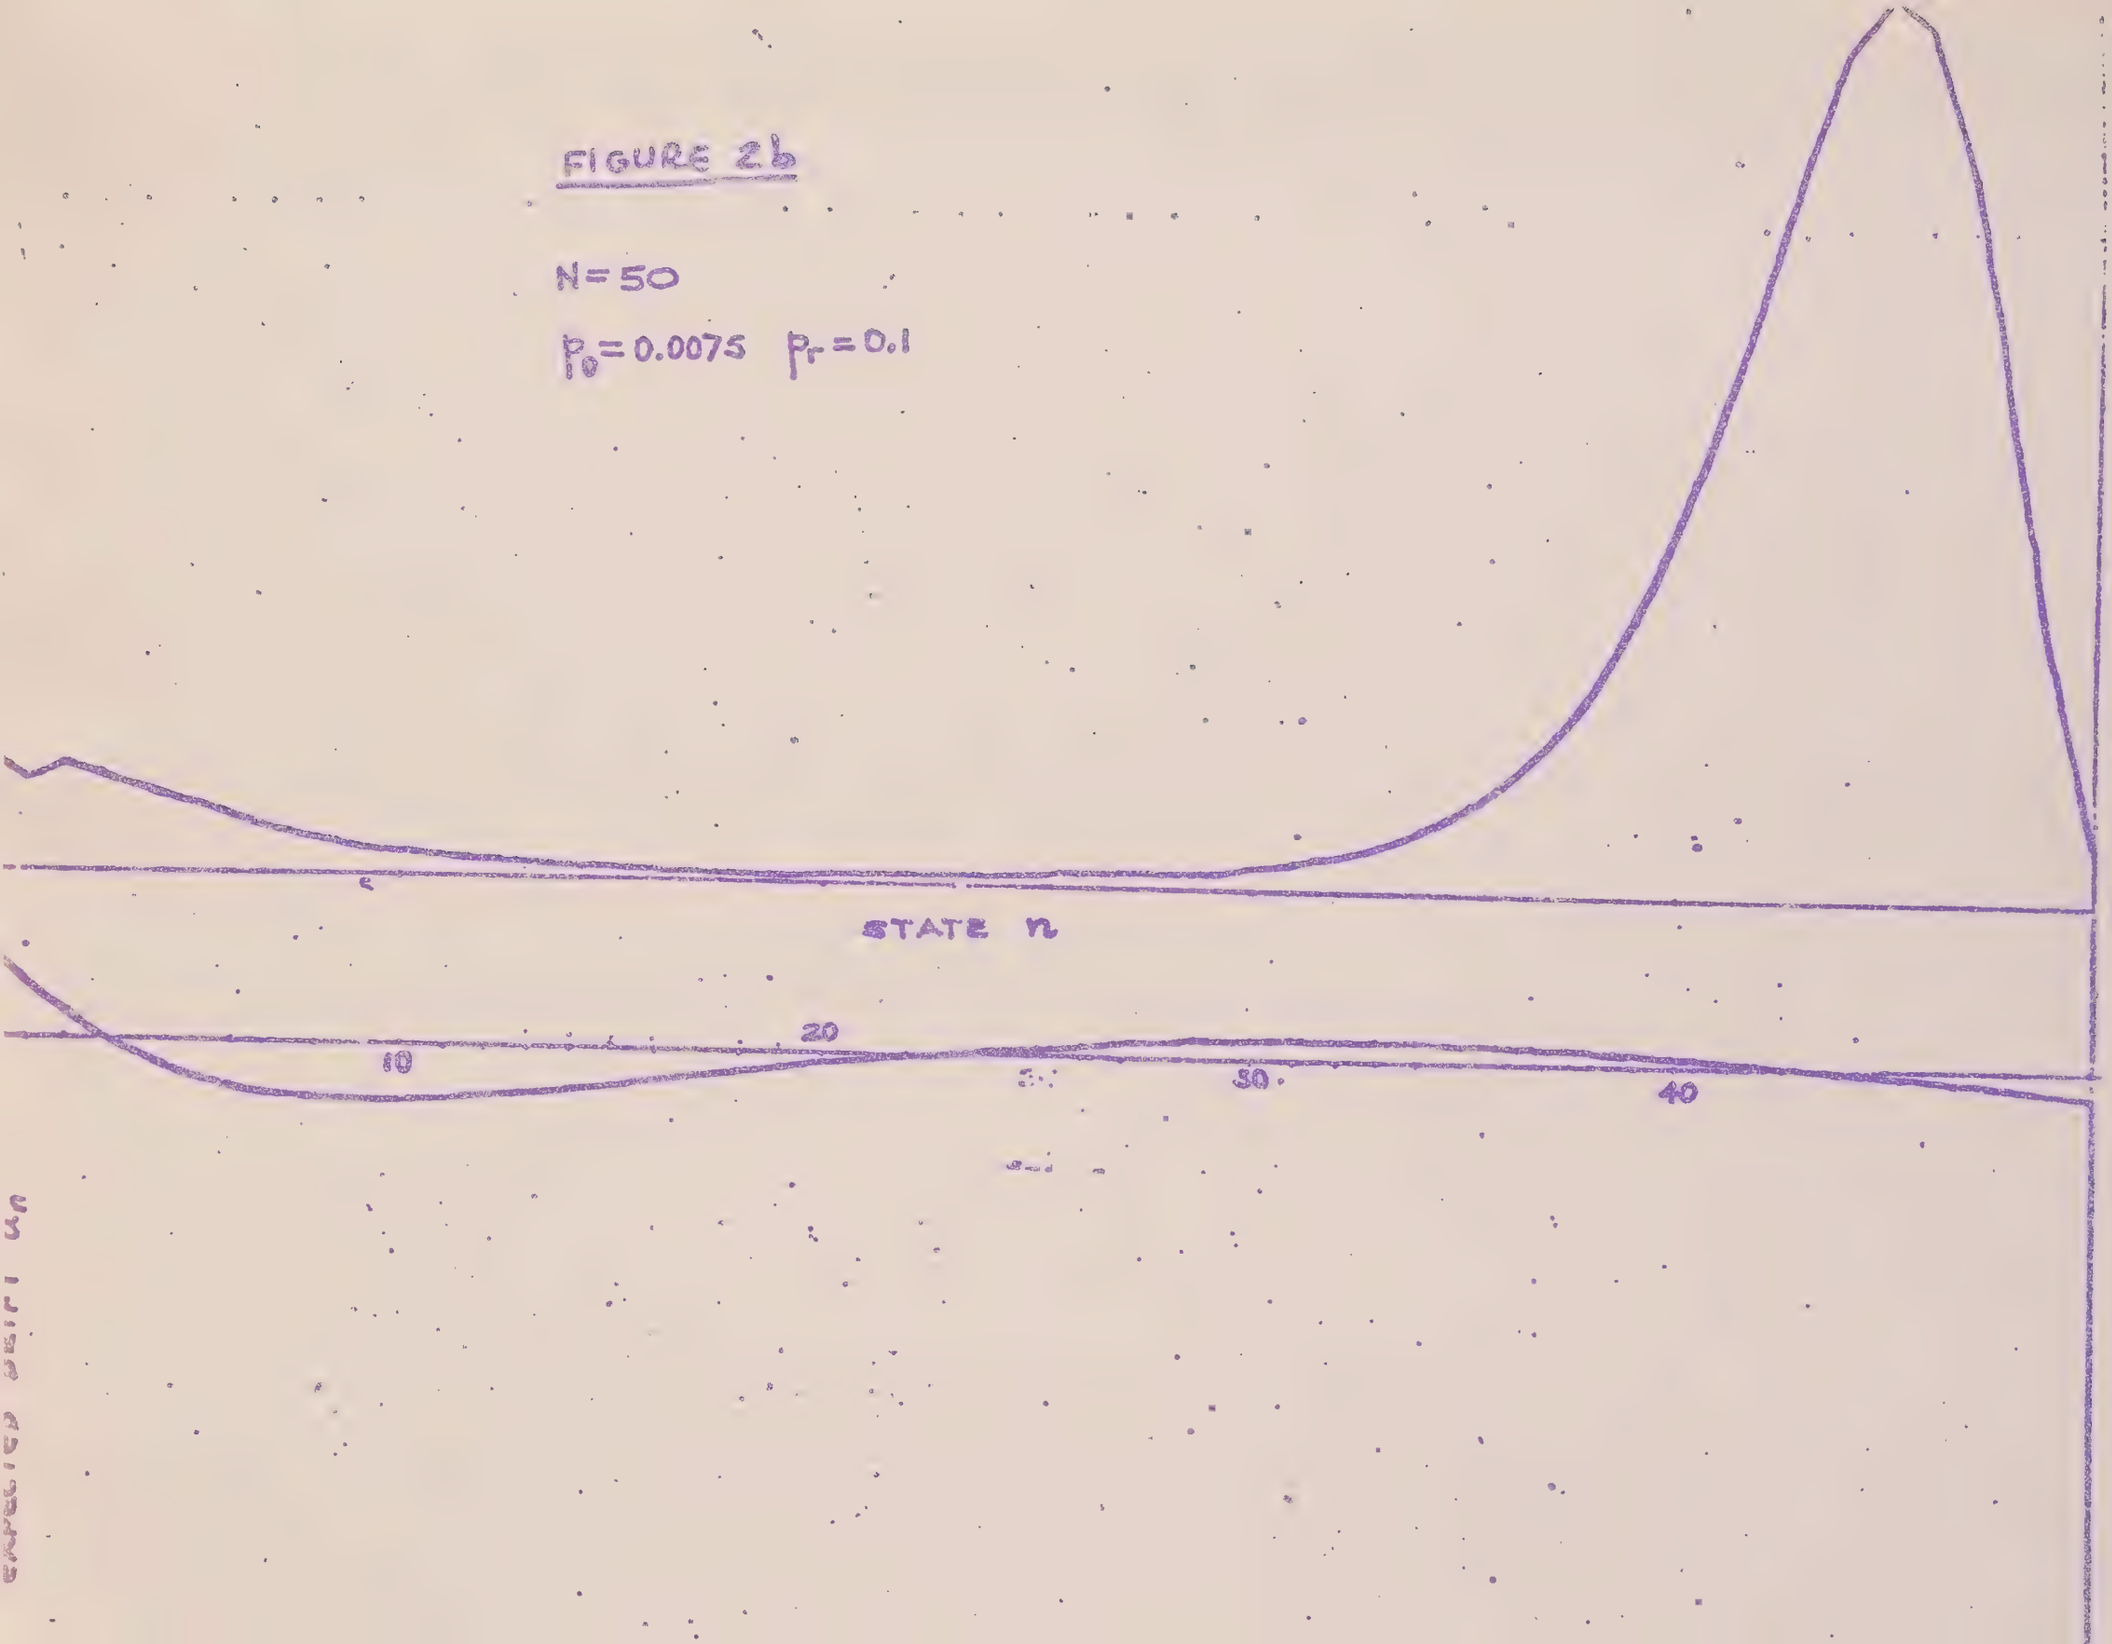

FIGURE 2c

$N = 50$

$P_0 = 0.01$ $P_r = 0.1$

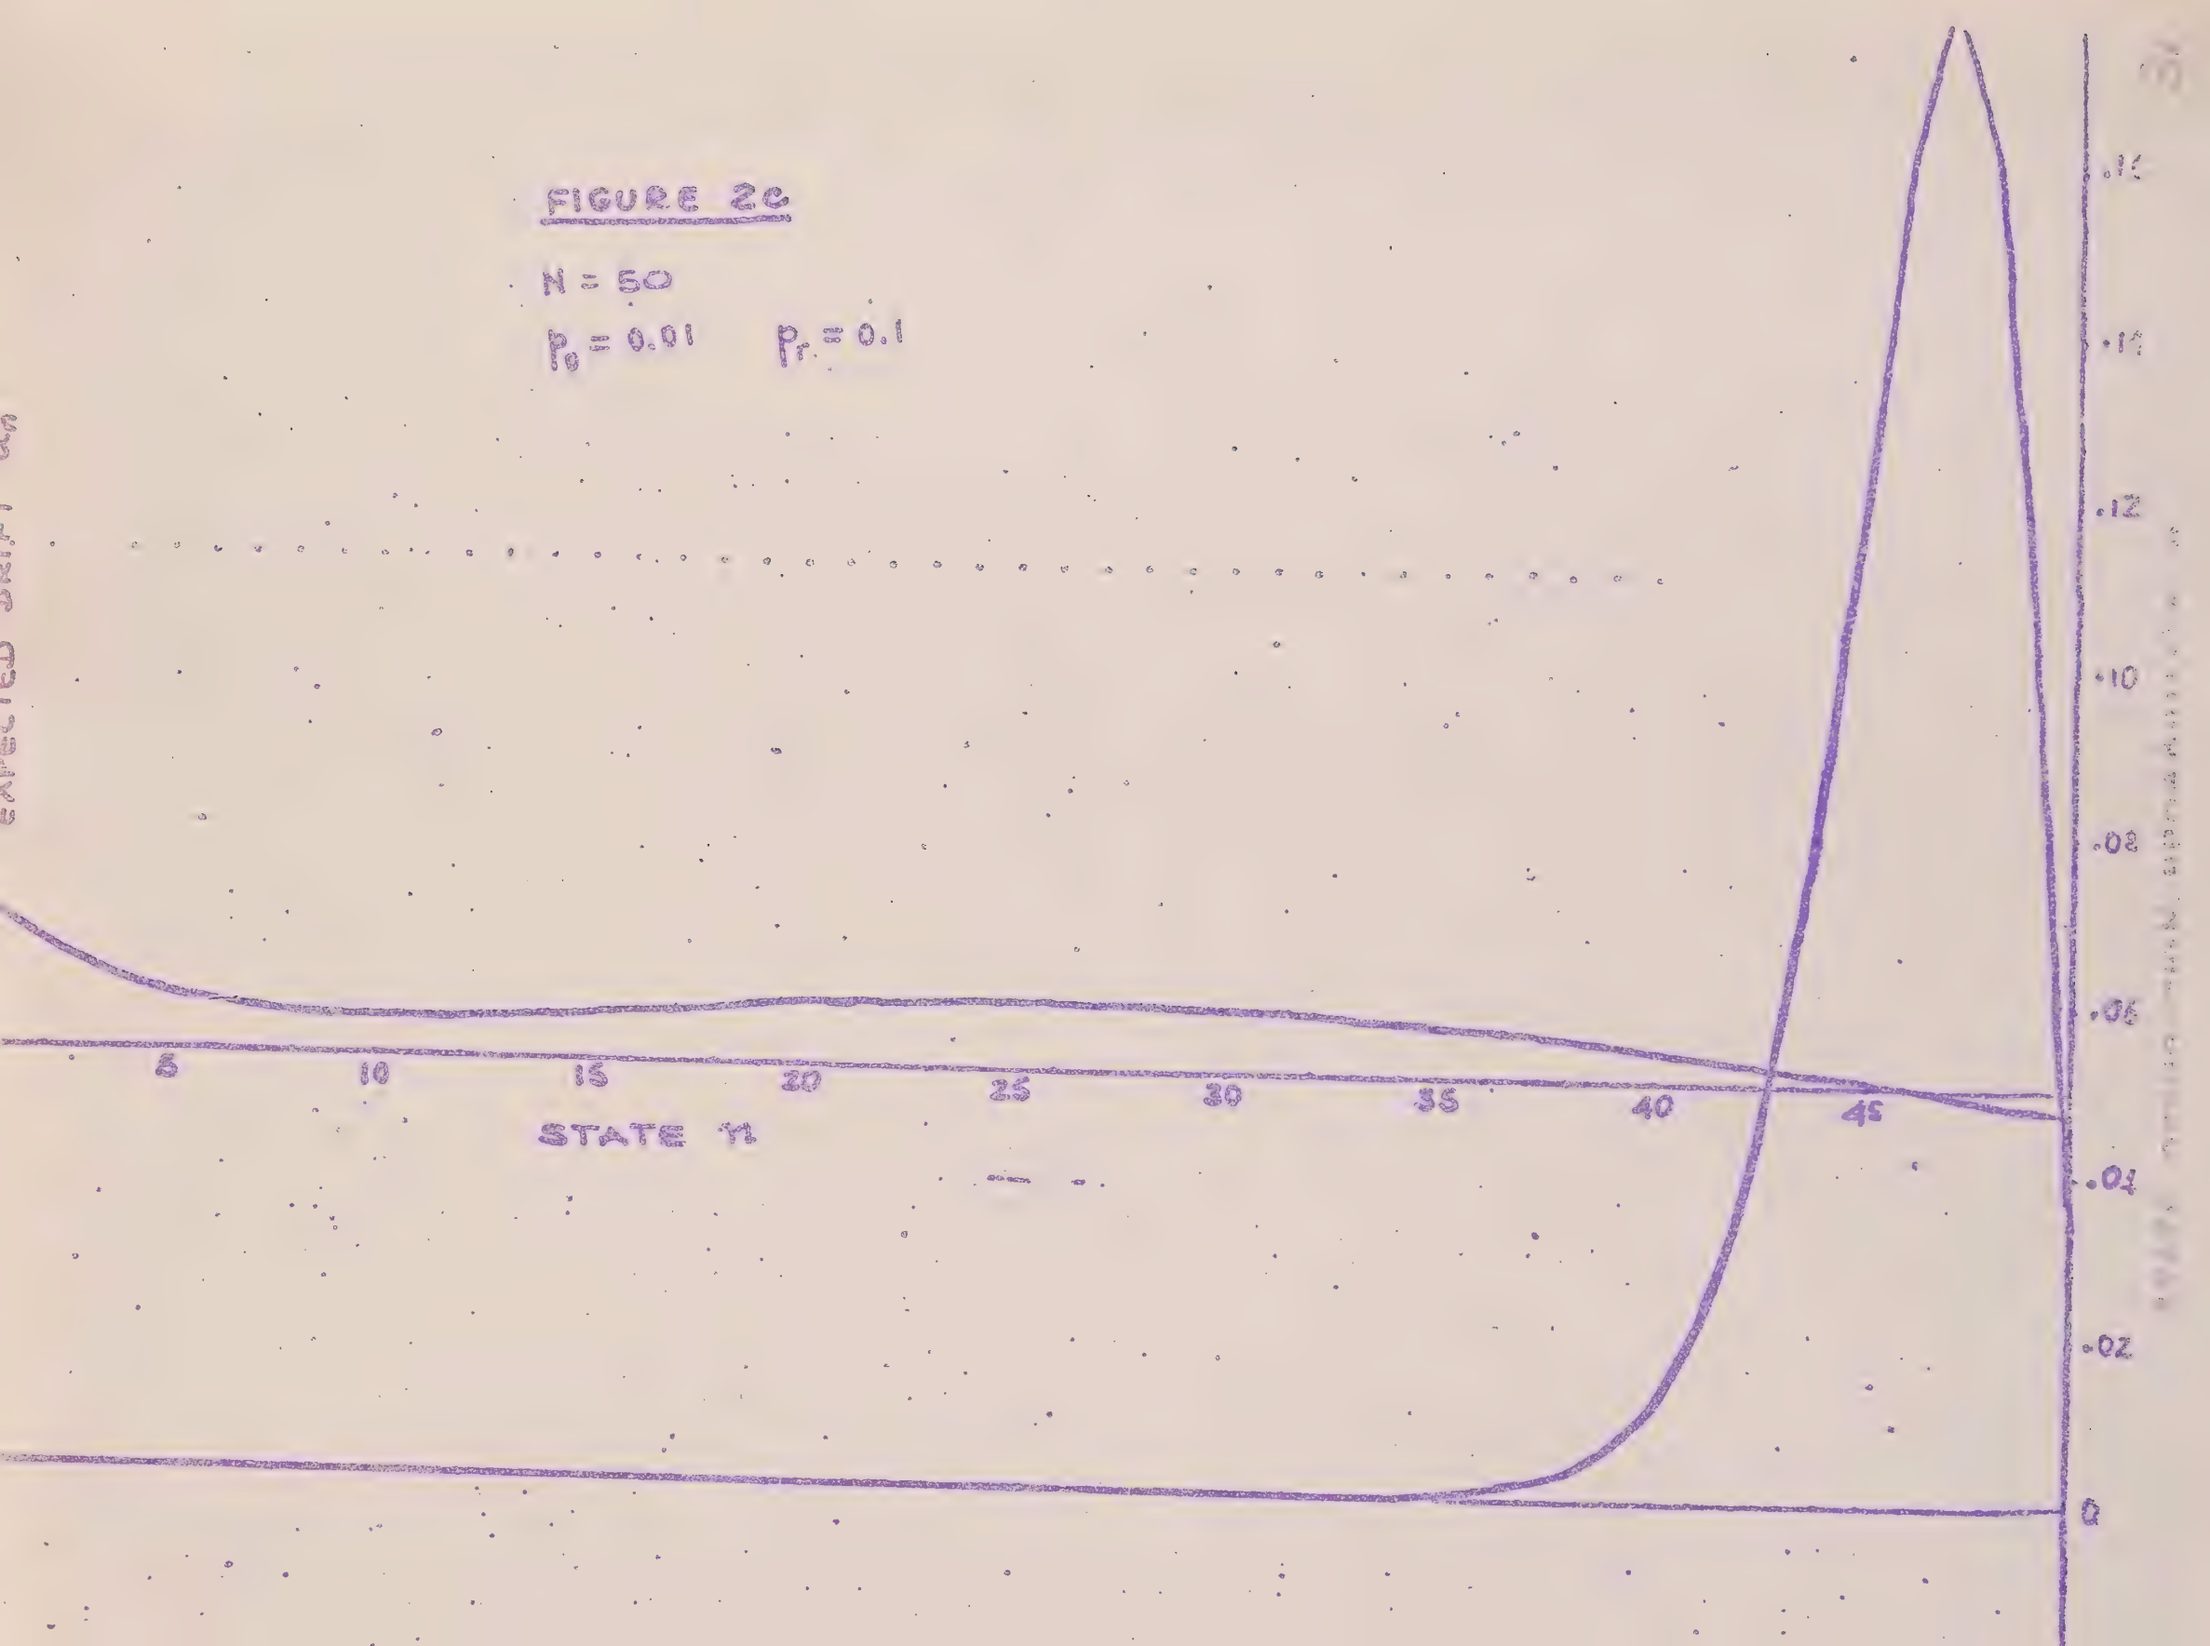

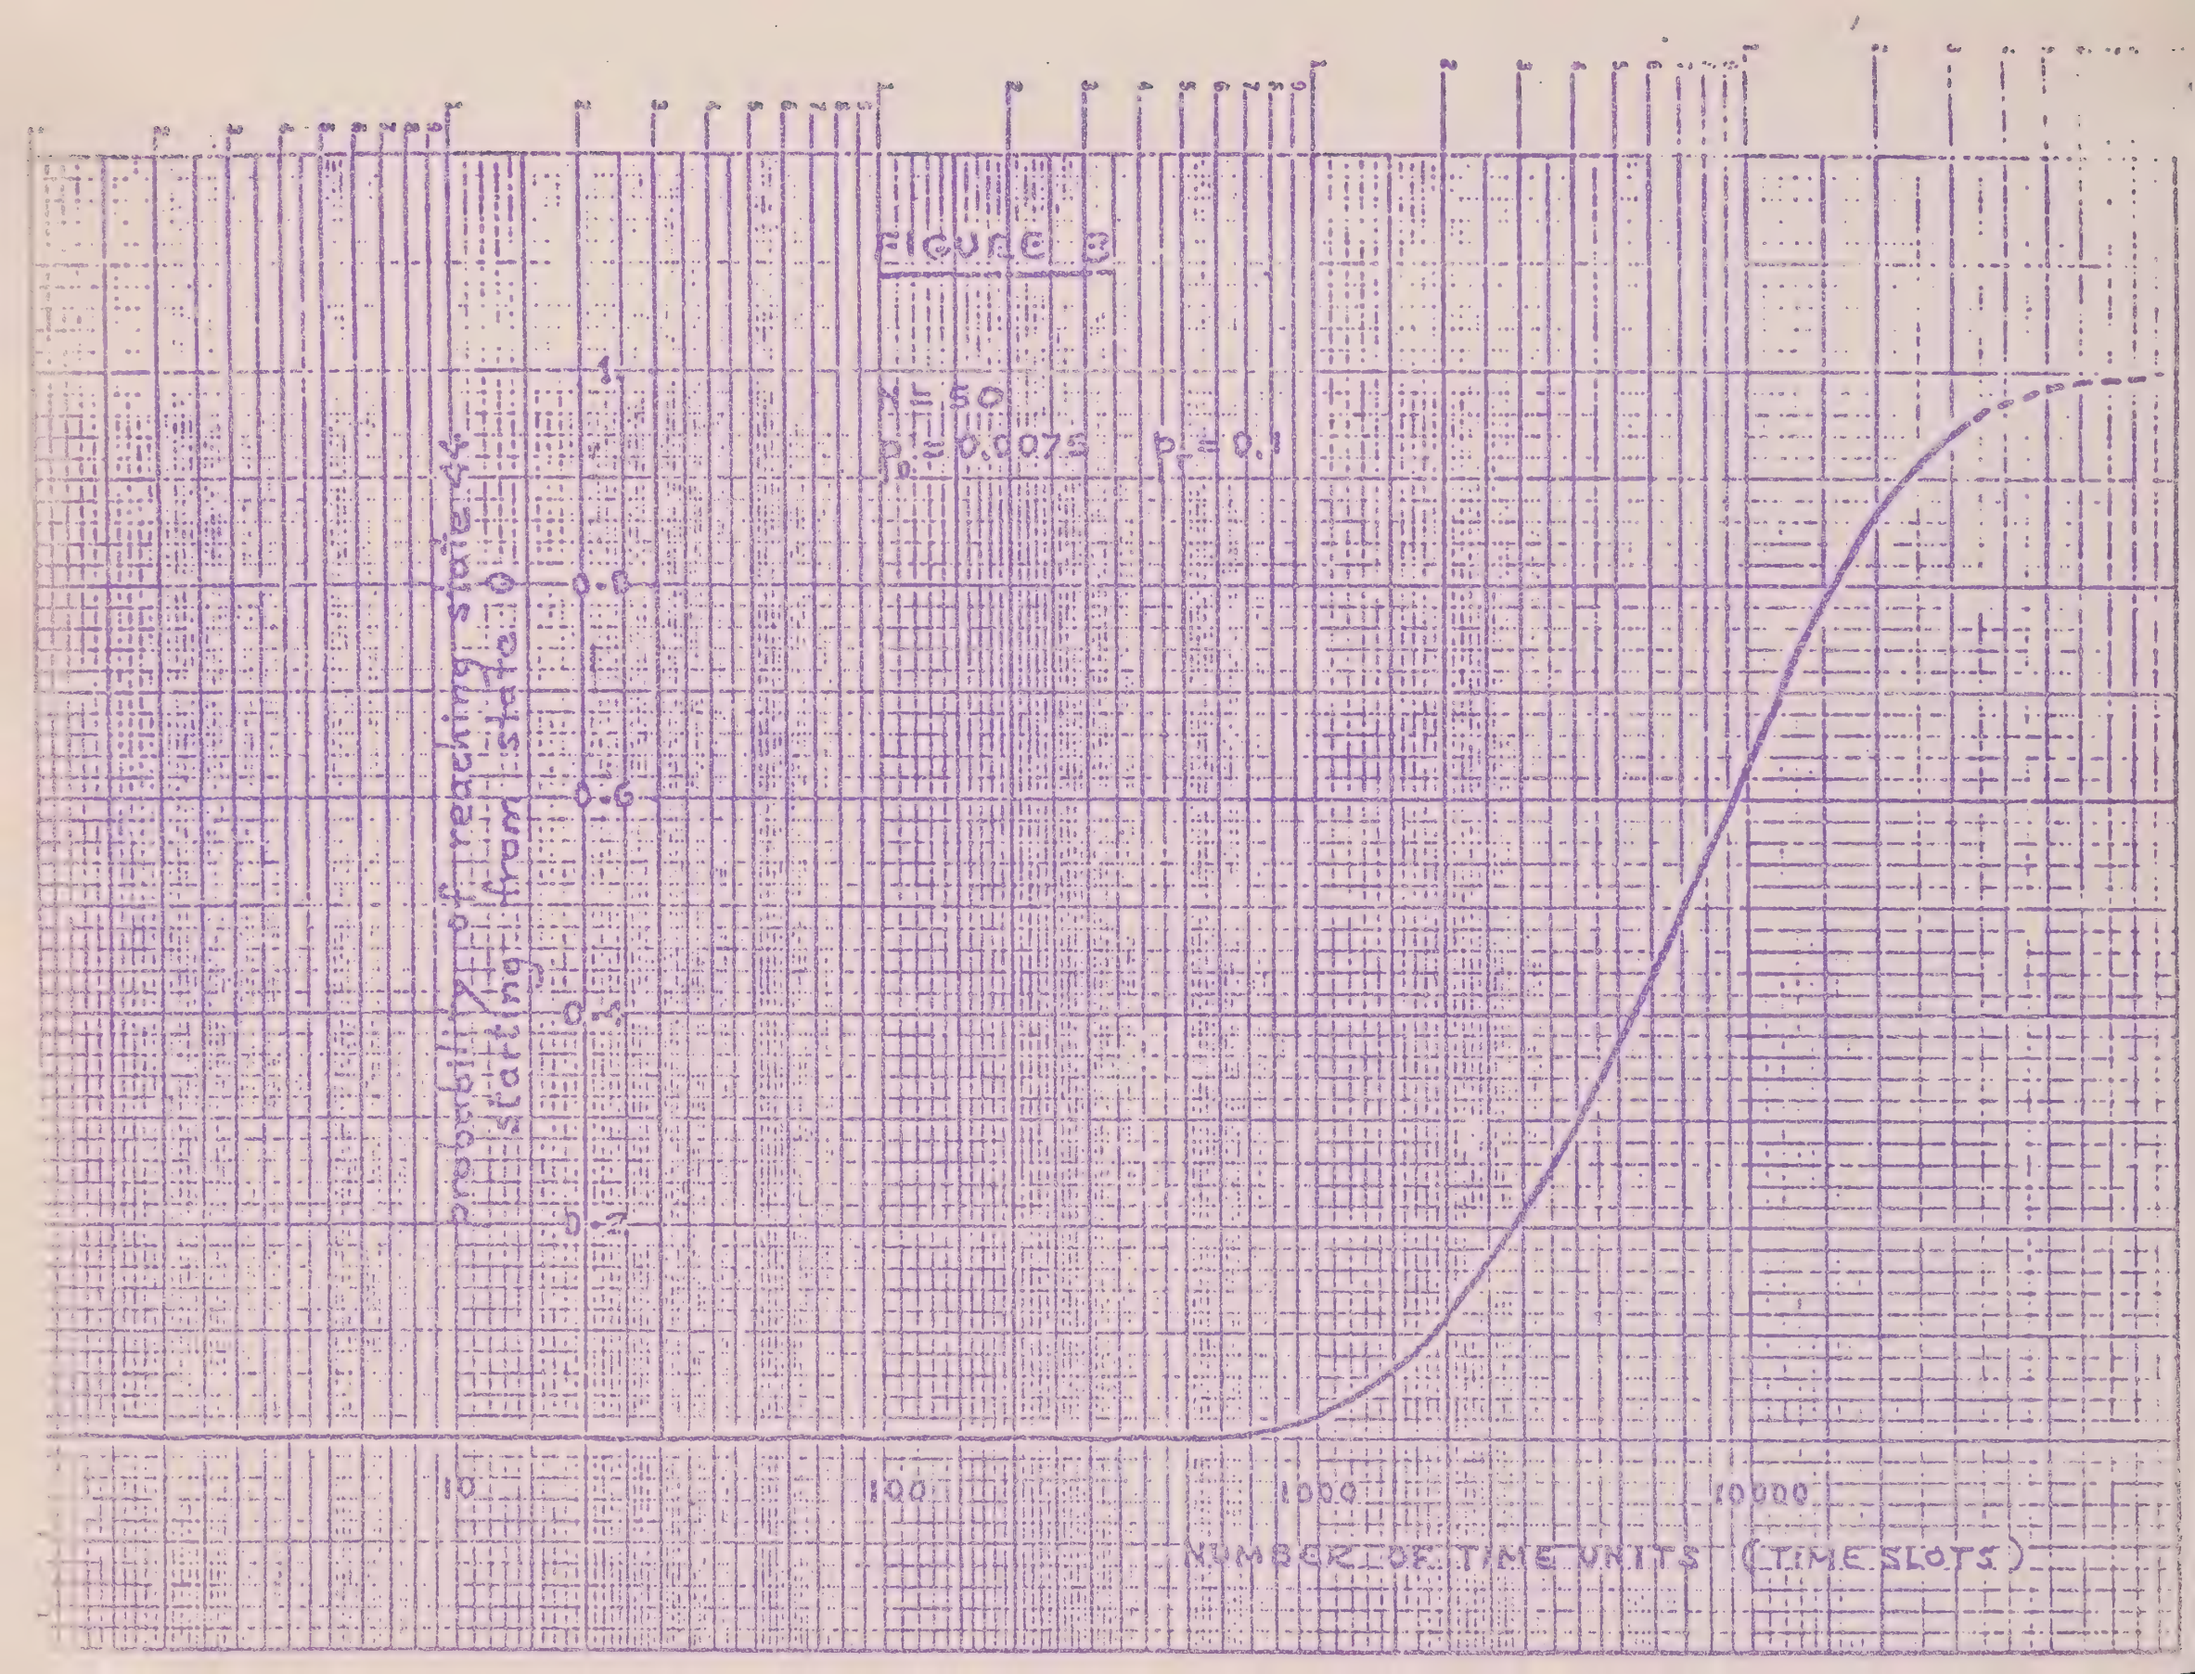
FIGURE 3
N = 50
$p_0 = 0.0075$
$p_1 = 0.1$
probability of rebeaching state 0 starting from state 0
0.8
0.6
0.4
0.2
0
100
1000
10000
NUMBER OF TIME UNITS (TIME SLOTS)

FIGURE 4

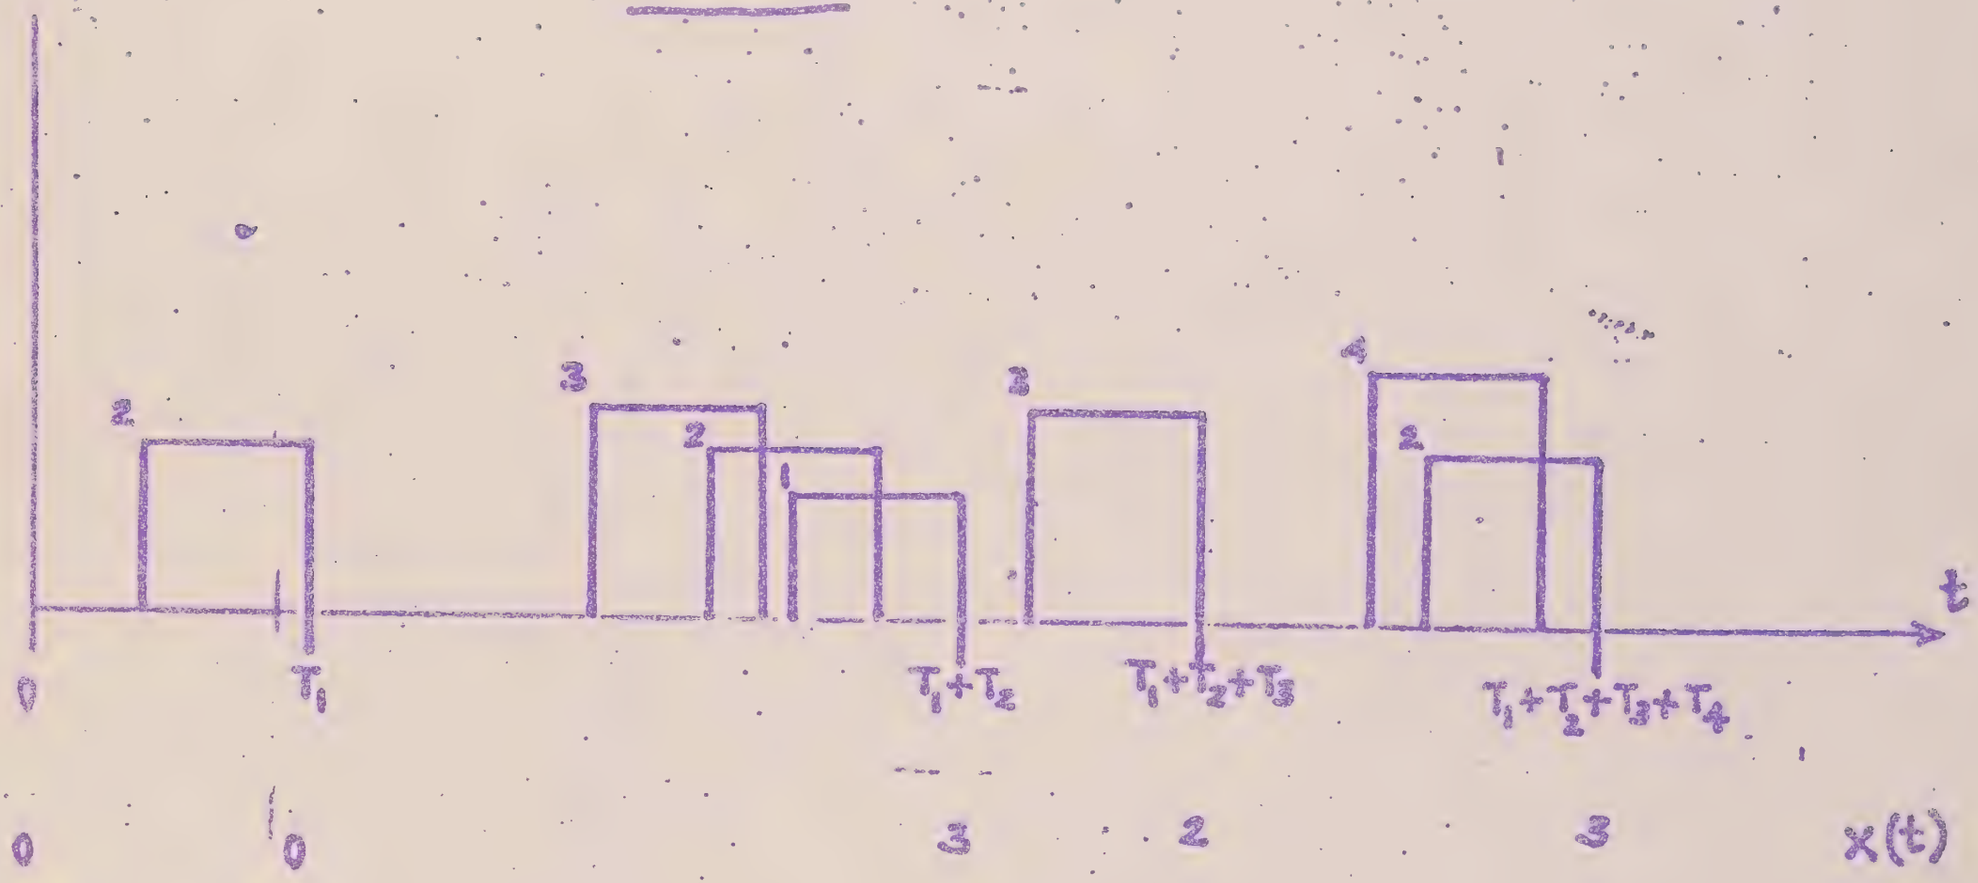

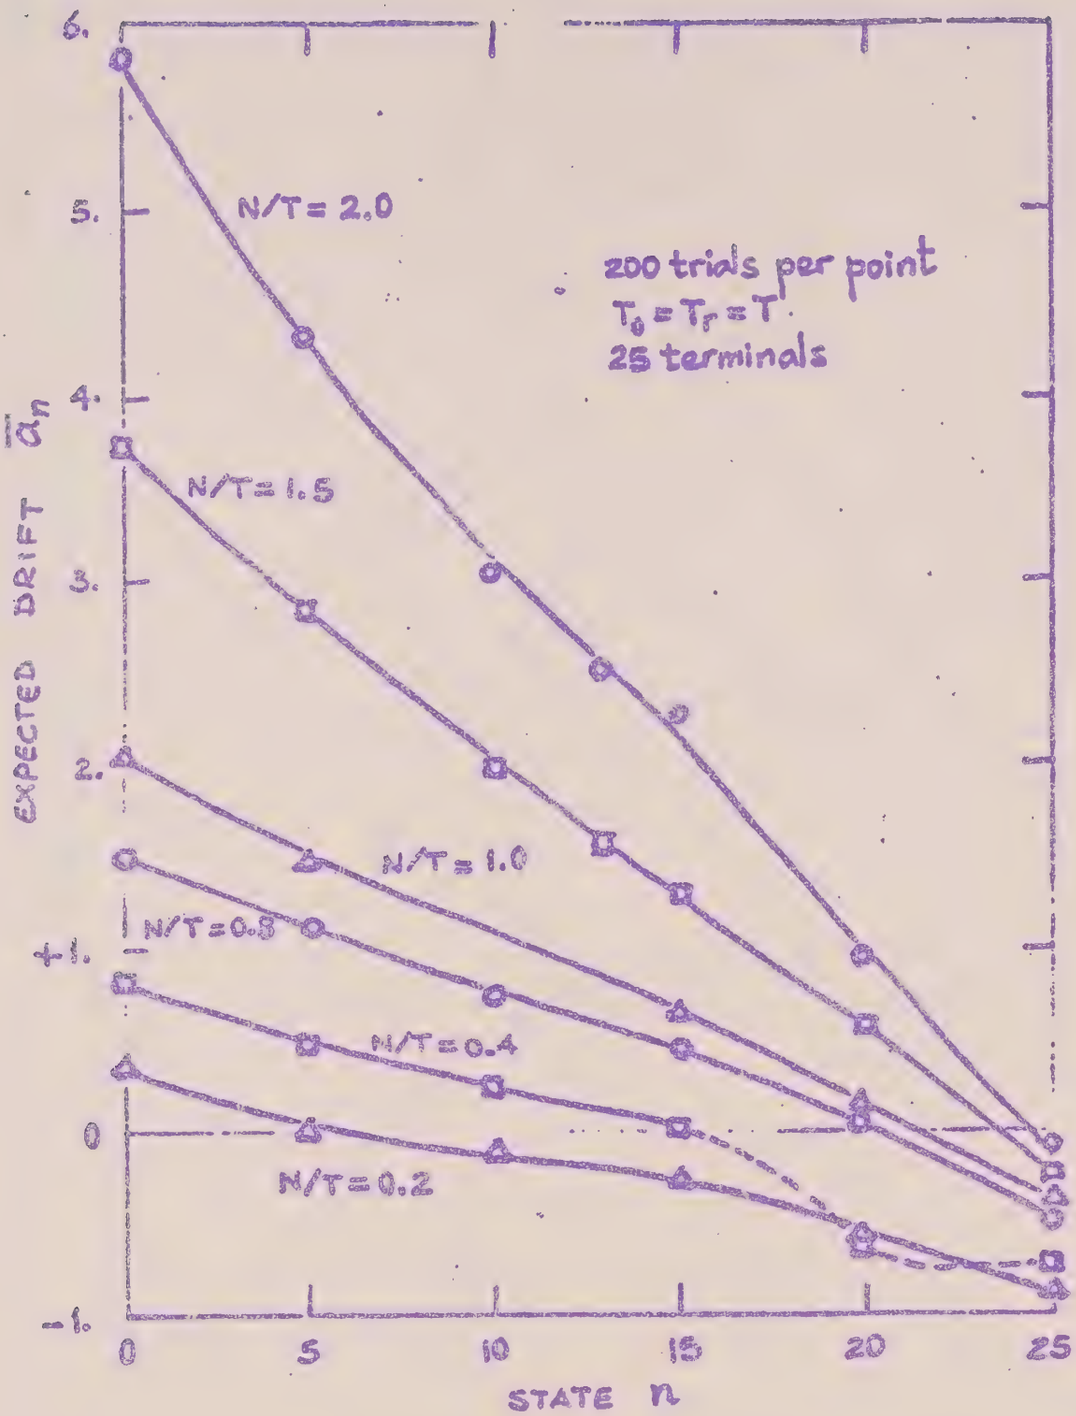
6.
5.
4.
3.
2.
+1.
0
-1.
EXPECTED DRIFT $\bar{a}_n$
N/T = 2.0
N/T = 1.5
N/T = 1.0
N/T = 0.8
N/T = 0.4
N/T = 0.2
200 trials per point
$T_0 = T_f = T$
25 terminals
0
5
10
15
20
25
STATE n

FIGURE 5

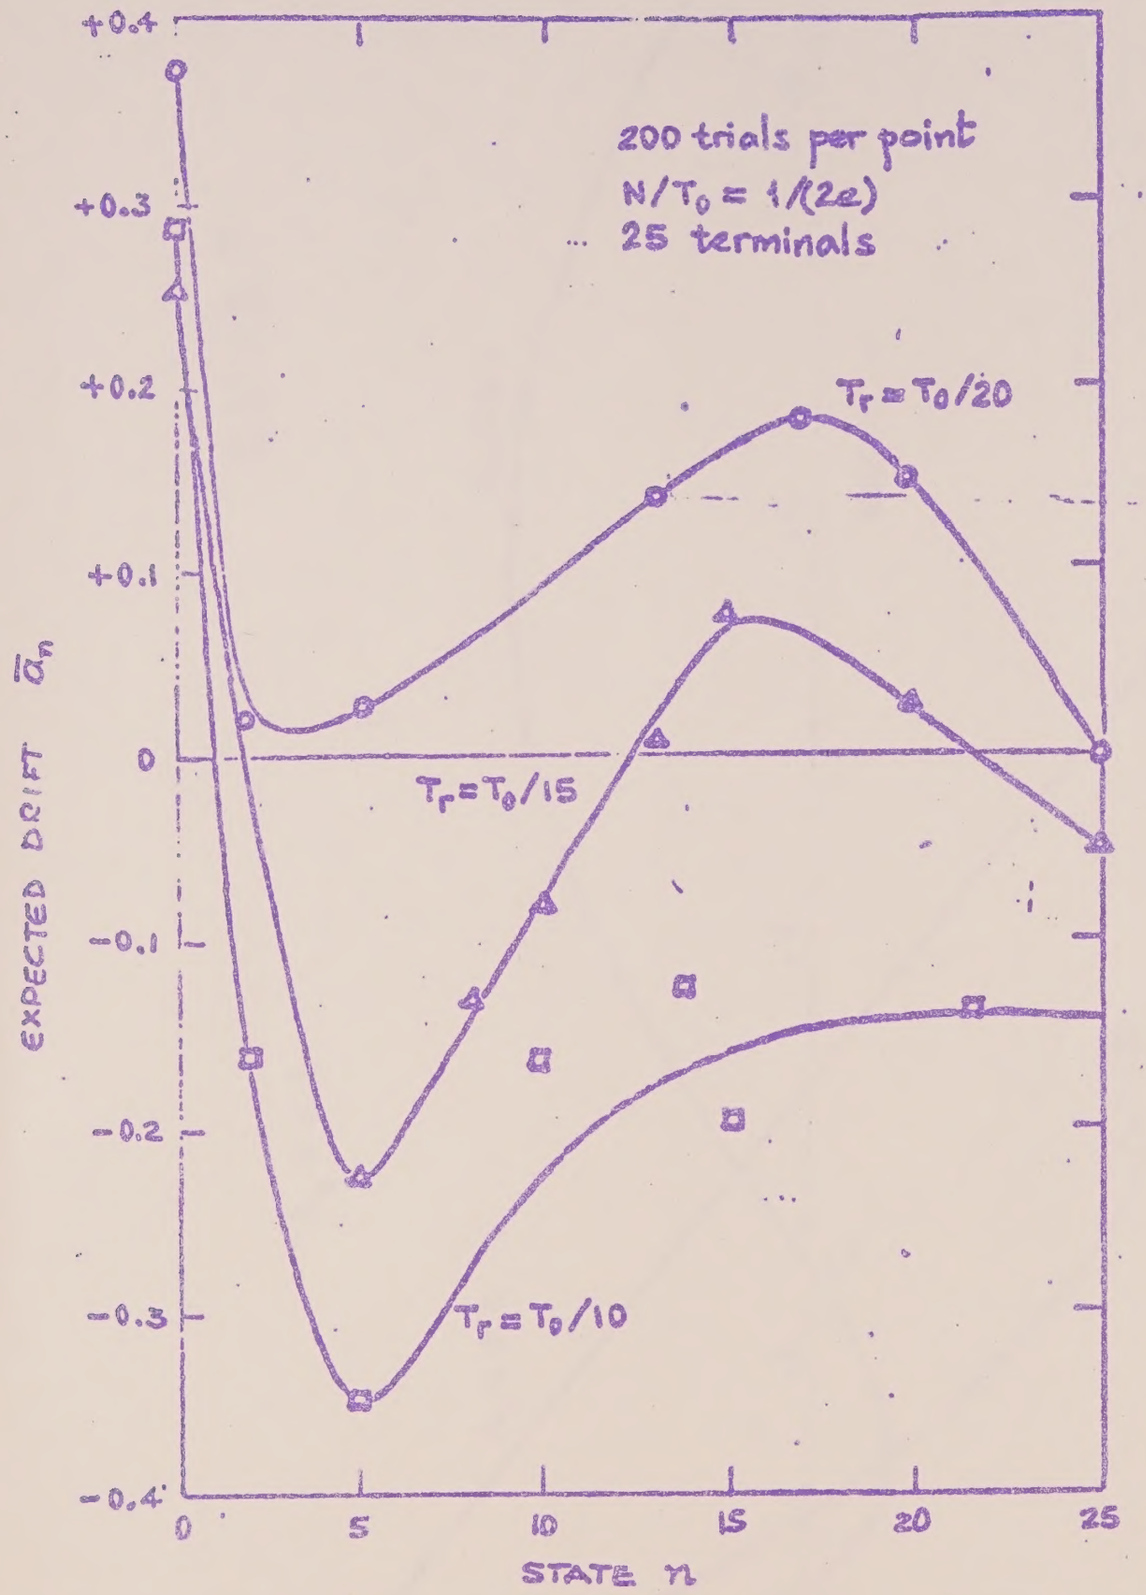
200 trials per point
N/T0 = 1/(2e)
25 terminals
Tr = T0/20
Tr = T0/15
Tr = T0/10
EXPECTED DRIFT ān
STATE n
+0.4
+0.3
+0.2
+0.1
0
-0.1
-0.2
-0.3
-0.4
0
5
10
15
20
25

FIGURE 6.

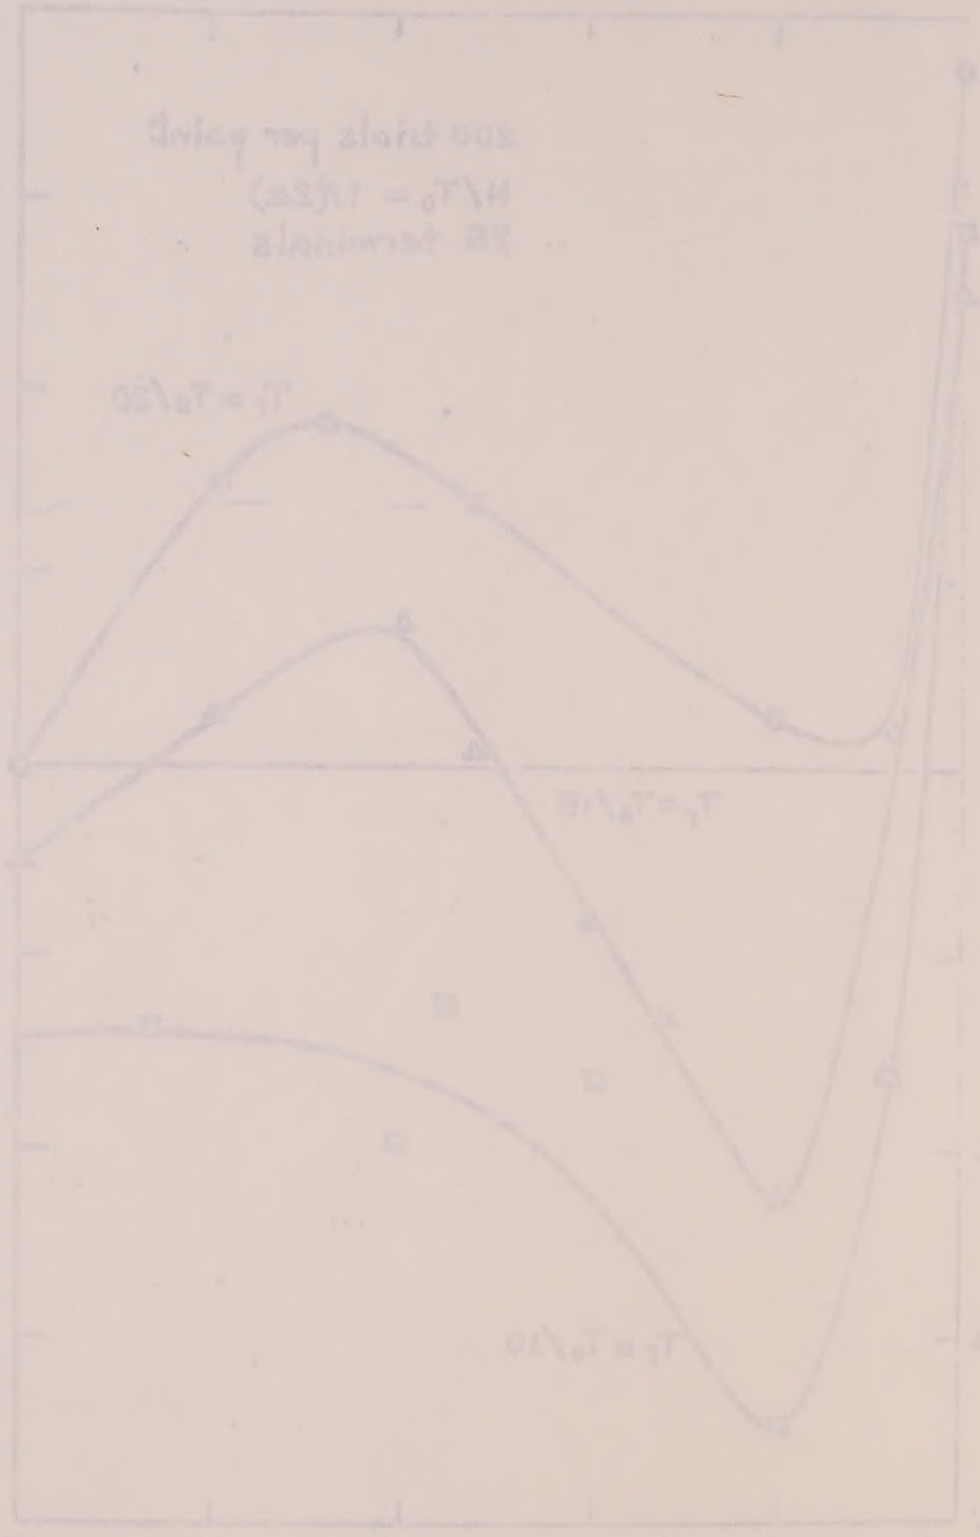

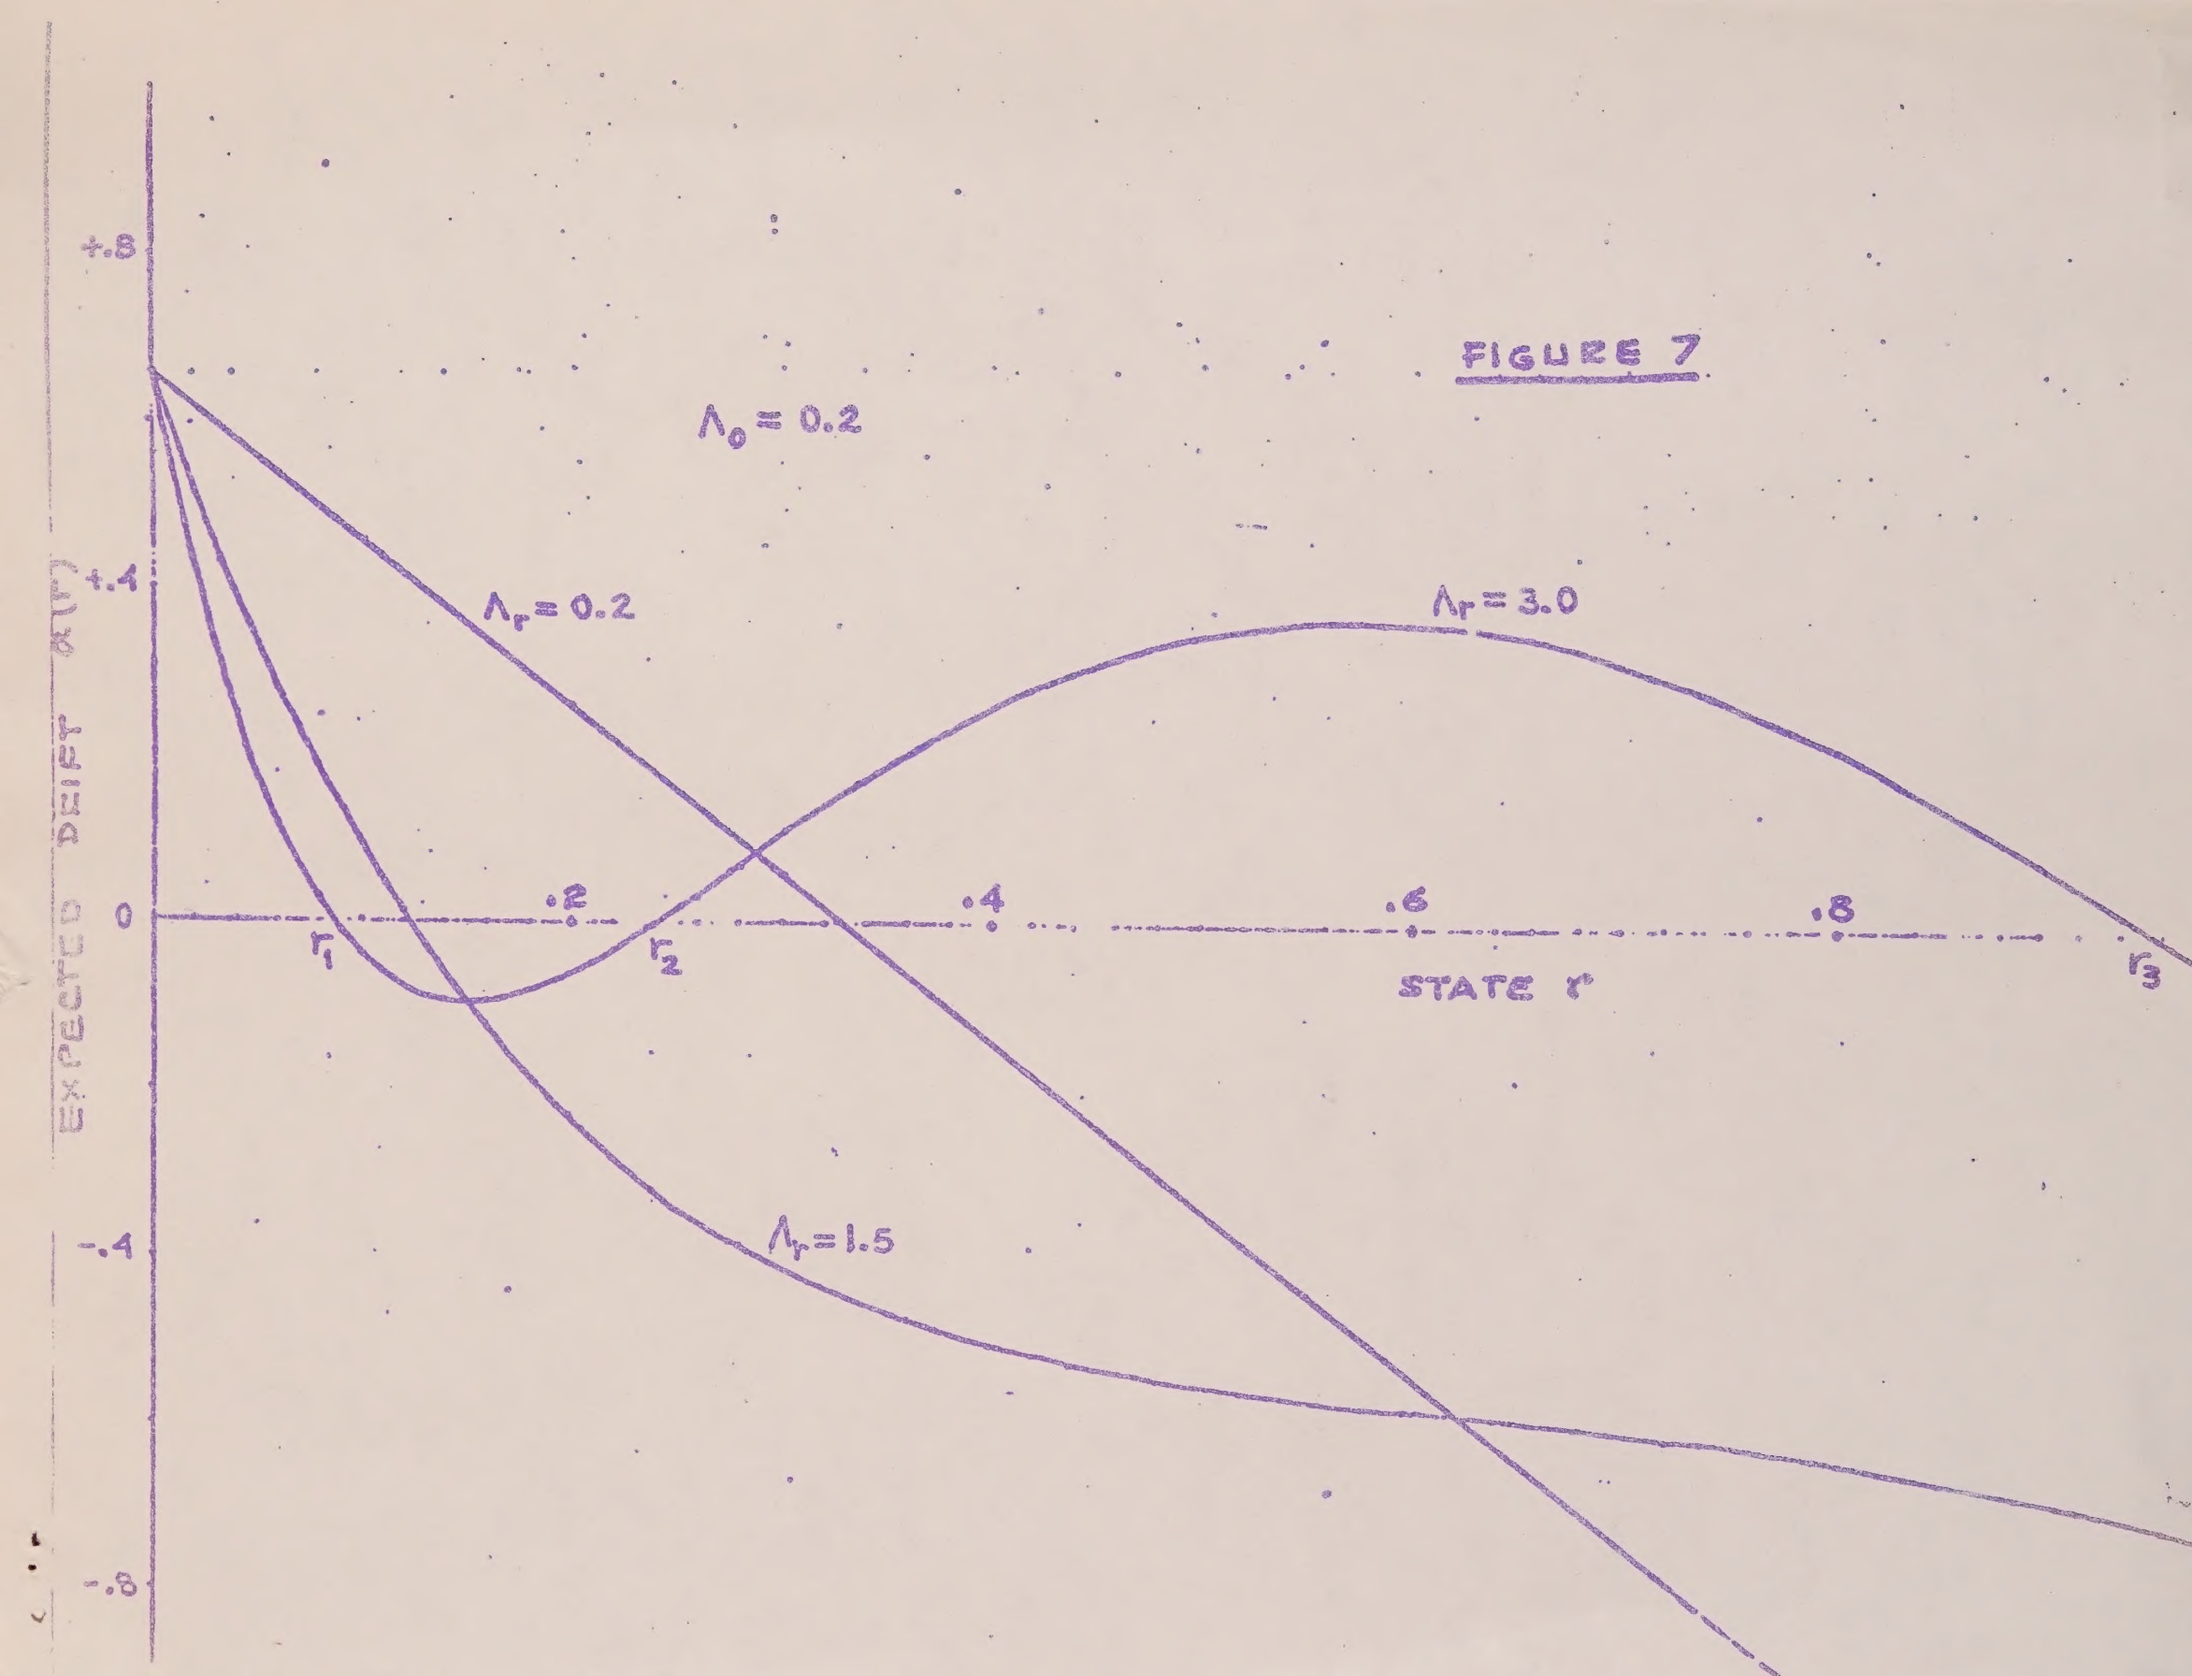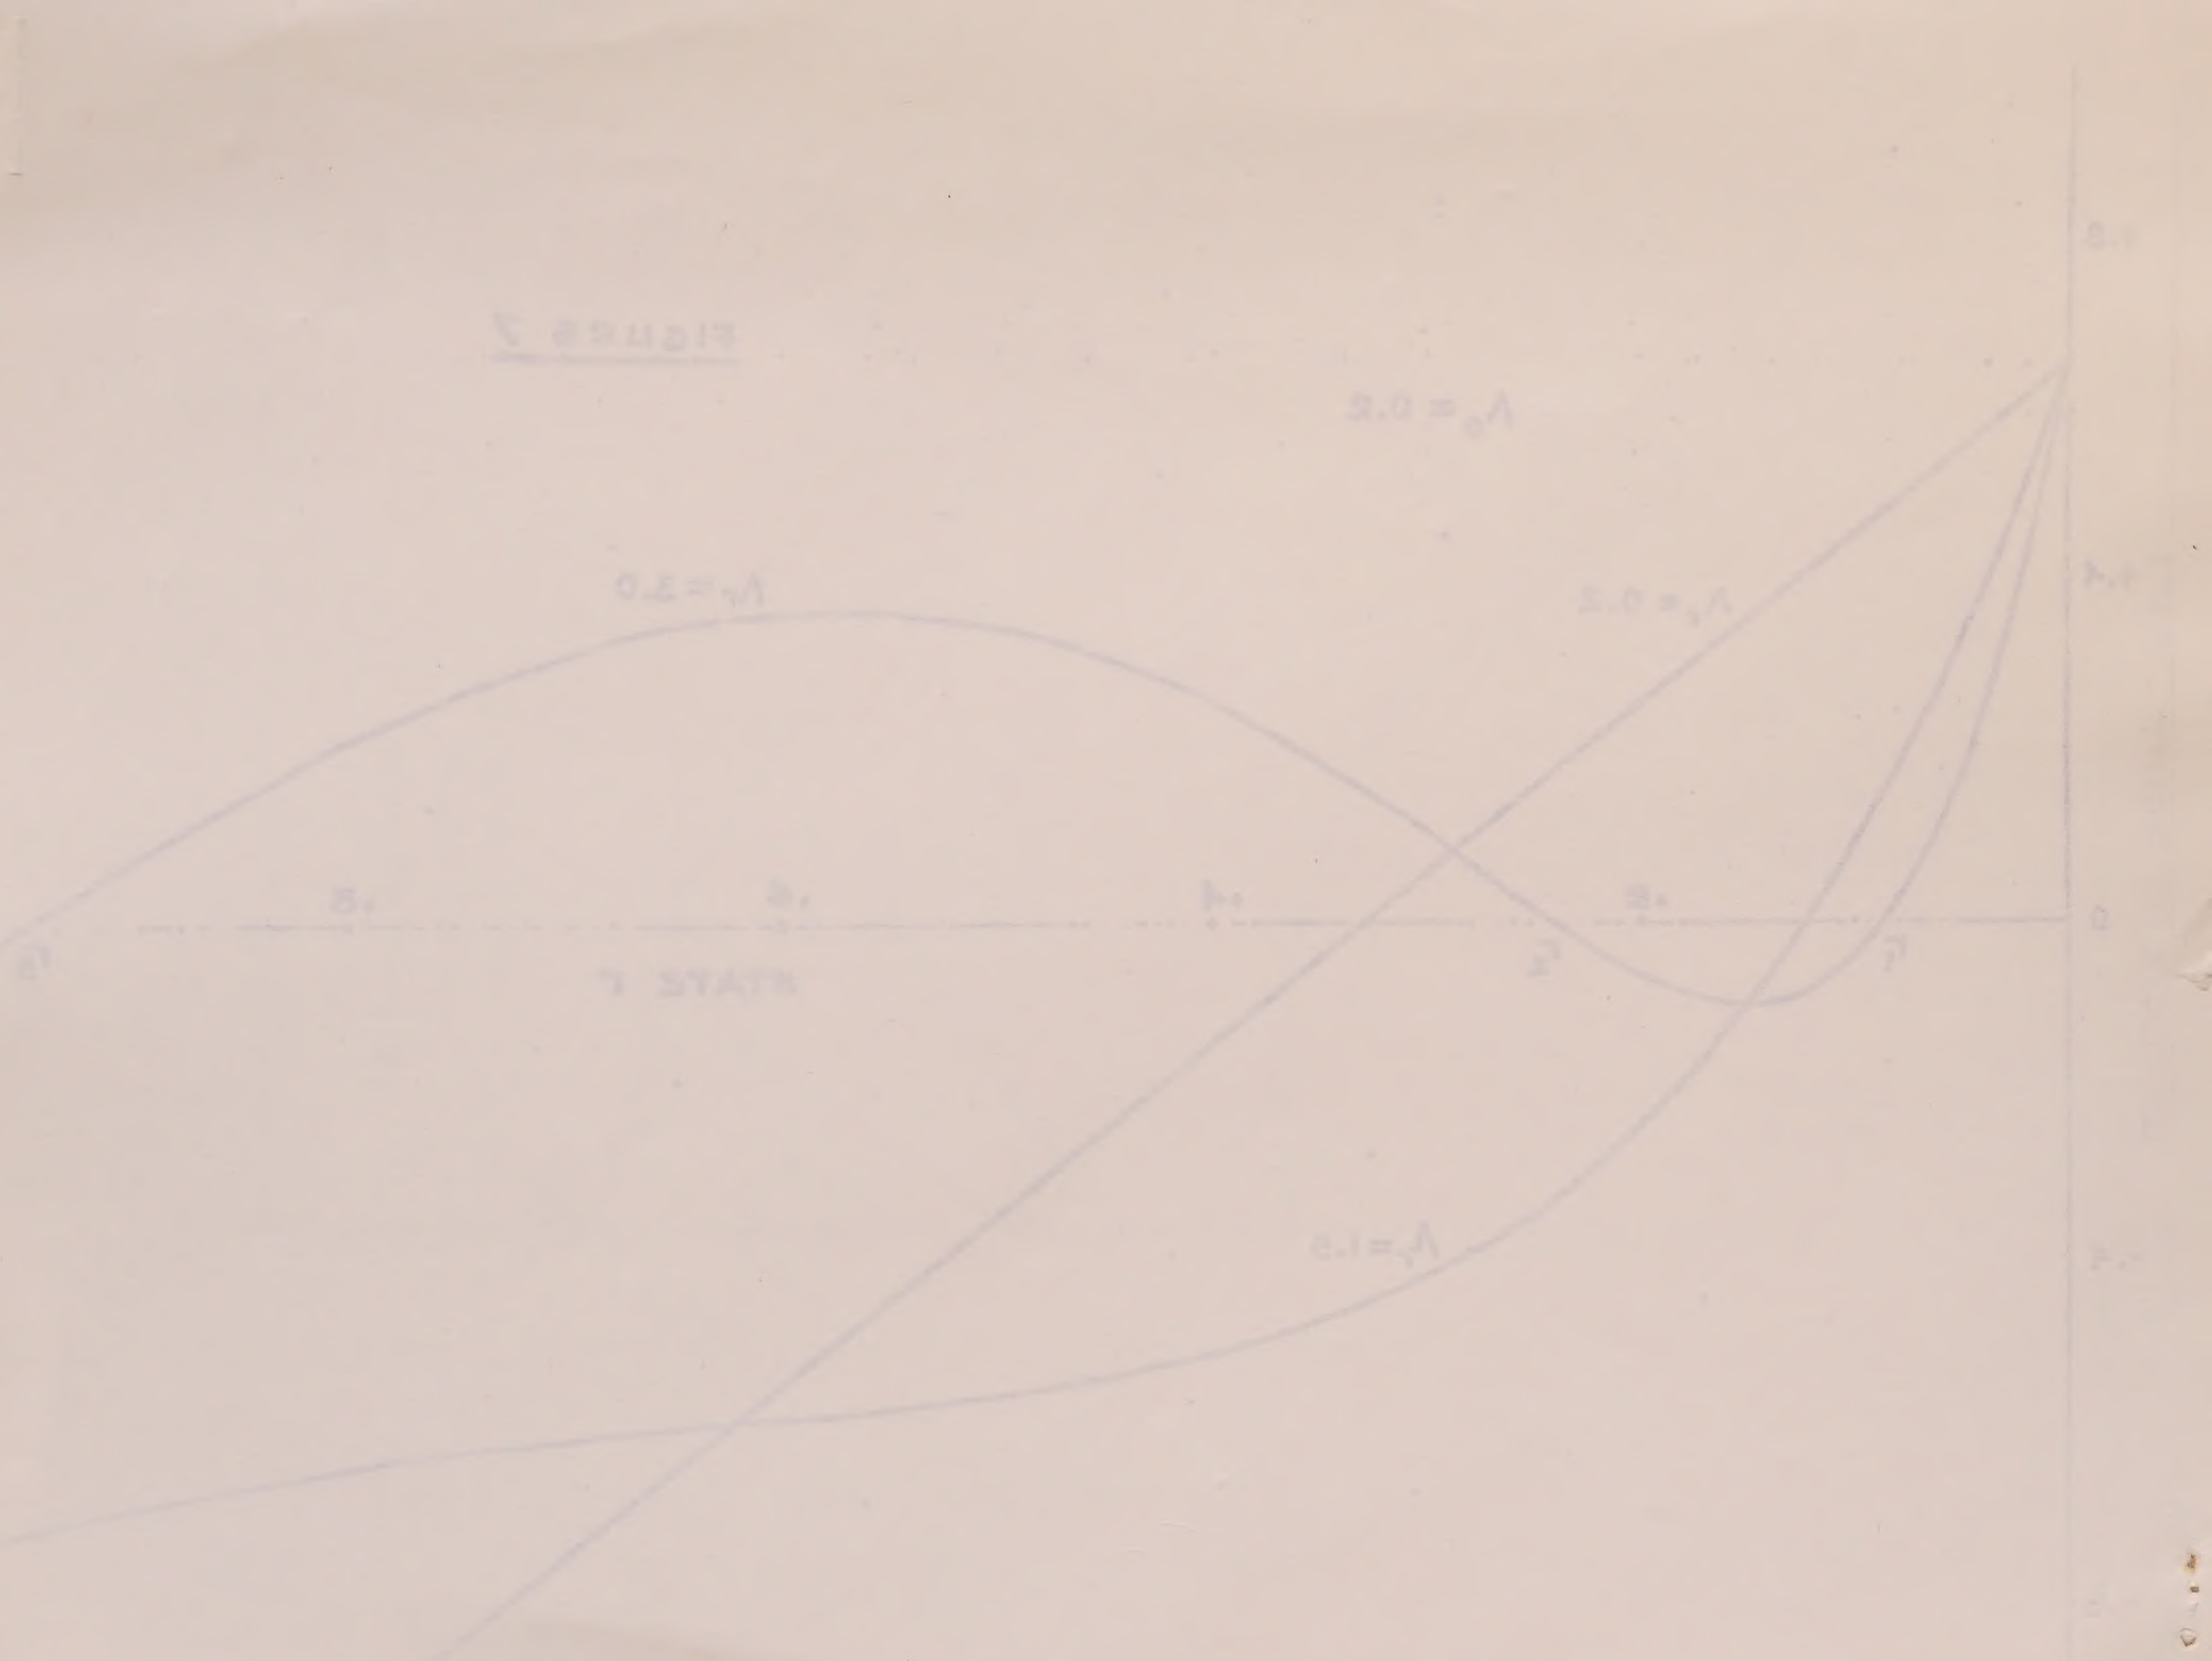
FIGURE 7
$\Lambda_0 = 0.2$
$\Lambda_r = 0.2$
$\Lambda_r = 3.0$
$\Lambda_r = 1.5$
+.8
+.4
0
-.4
-.8
.2
.4
.6
.8
r_1
r_2
r_3
STATE r'
EXPECTED DRIFT $\delta(r')$